

# **How Are We Doing?**

## **A Self-Assessment of the Quality of Services and Systems at NERSC (2000)**

Issued April 2001 by the  
High Performance Computing Department  
William T. Kramer, Department Head

National Energy Research Scientific Computing Center  
Ernest Orlando Lawrence Berkeley National Laboratory

Berkeley Lab Report Number LBNL-47712

This work was supported by the Department of Energy's Office of Science, Office of Advanced Scientific Computing Research, Mathematical, Information and Computational Sciences Division, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

## **DISCLAIMER**

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

## **ACKNOWLEDGMENT**

This work was supported by the Department of Energy's Office of Science, Office of Advanced Scientific Computing Research, Mathematical, Information and Computational Sciences Division, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

This report was compiled by Jon Bashor. For more information, call 510-486-5849, or write to JBashor@lbl.gov. To learn more about the National Energy Research Scientific Computing Center, visit our web site at: <<http://www.nersc.gov>>.

**CONTENTS**

INTRODUCTION — HIGHLIGHTS OF THE PAST YEAR..... 1

RELIABLE AND TIMELY SERVICE

    Goal: Have all systems and support functions provide reliable and timely service to their clients. .... 6

INNOVATIVE ASSISTANCE

    Goal: NERSC aims to provide its clients with new ideas, new techniques and new solutions to their scientific computing issues. .... 11

TIMELY AND ACCURATE INFORMATION

    Goal: Provide timely and accurate information and notification of system changes to the client community so they can most effectively use the NERSC systems. .... 14

NEW TECHNOLOGIES, EQUIPMENT, SOFTWARE, AND METHODS

    Goal: Ensure that future high-performance technologies are available to Office of Science computational scientists. .... 16

WISE TECHNOLOGY INTEGRATION

    Goal: Ensure all new technology and changes improve (or at least do not diminish) service to our clients. .... 18

PROGRESS MEASUREMENT

    Goal: As a national facility serving thousands of researchers, NERSC has a responsibility to our clients to measure and report on how we’re doing in terms of providing service, support and facilities. .... 21

HIGH-PERFORMANCE COMPUTING CENTER LEADERSHIP

    Goal: Improve methods of managing systems within NERSC and be the leader in large-scale computing center management. .... 25

TECHNOLOGY TRANSFER

    Goal: Export knowledge, experience, and technology developed at NERSC, particularly to and within NERSC client sites. .... 29

STAFF EFFECTIVENESS

    Goal: Leverage staff expertise and capabilities to increase efficiency and effectiveness. .... 48

PROTECTED INFRASTRUCTURE

    Goal: Provide a secure computing environment for NERSC sponsors. .... 51

CONCLUSION..... 52



**EXECUTIVE SUMMARY**

This fourth annual self-assessment of the systems and services provided by the U.S. Department of Energy's National Energy Research Scientific Computing Center describes the efforts of the NERSC staff to support advanced computing for scientific discovery. Our staff applies experience and expertise to provide world-class systems and unparalleled services for NERSC users. At the same time, members of our organization are leading contributors to advancing the field of high-performance computing through conference presentations, published papers, collaborations with scientific researchers and through regular meetings with members of similar institutions. We believe that, by any measure, the results of our efforts underscore NERSC's position as a global leader in scientific computing.



## **INTRODUCTION — HIGHLIGHTS OF THE PAST YEAR**

In the year 2000, 25 years after it went on line with a machine of its own and began building a legacy of HPC contributions, NERSC made a series of moves to clearly define itself as a leader in 21<sup>st</sup> century computational science. These steps included the initial installation of the world's most powerful unclassified supercomputer, the move to a state-of-the art computing facility in Oakland, and the purchase of a 160-processor IBM cluster to be evaluated for its suitability in meeting the production demands of NERSC users.

NERSC also made news on the strength of work by staff and users in tapping the center's expertise and facilities. These included data analysis which shaped findings about the geometry of the universe, a team of NERSC researchers chosen to have their work included in the Smithsonian Institution's permanent collection, and the release of a software library called Berkeley Lab AMR (Adaptive Mesh Refinement).

While making these advances, the NERSC staff continued to provide our user community with highly rated support services and with an unmatched availability of computing resources. Even during the move to Oakland in October, disruption to our users was minimal and our IBM SP remained in almost constant service.

To ensure that our staff remains focused on meeting the needs of NERSC and advancing computational science in supporting DOE's mission areas, we regularly establish a series of related goals and then assess our performance against them. As in the past three editions of our annual "How Are We Doing" report, we support our evaluation with statistics, anecdotes and documentation.

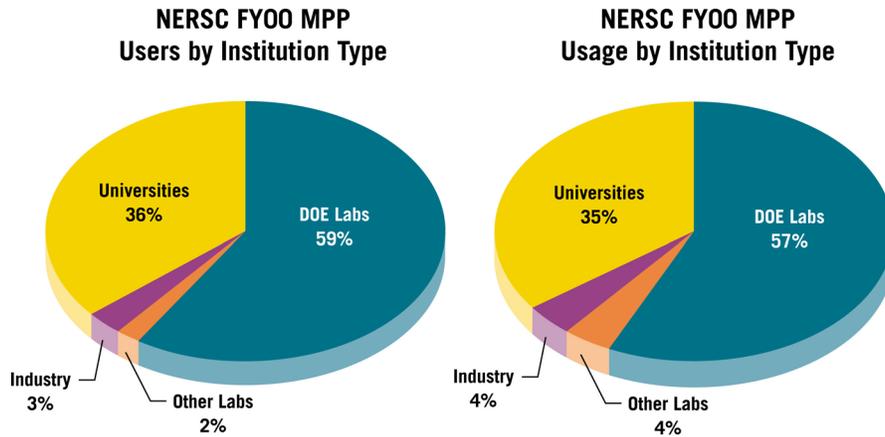
Meeting the diverse requirements of NERSC users is one of the drivers for constantly improving our systems and services. The following paragraphs and graphics illustrate the wide range of institutions, scientific disciplines and geographic distribution which characterize our user community.

### **A Statistical Snapshot of NERSC Users**

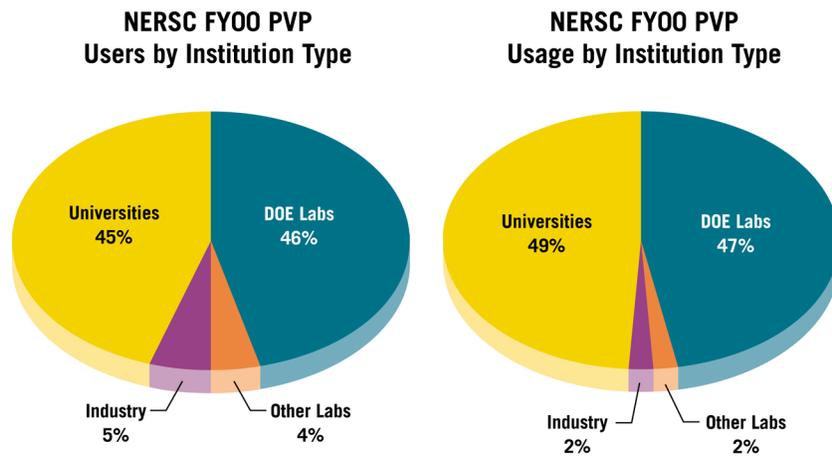
Meeting the computational science needs of the DOE Office of Science encompasses a broad range of researchers in terms of scientific disciplines, geographic location or home institution. Here are some statistics on the NERSC user community.

#### **By Type of Institution**

NERSC predominantly serves users at DOE national laboratories, who account for 59 percent of the center's MPP use. Universities account for 36 percent of MPP use, other labs claim 2 percent and use by industry is about 3 percent. By calculating MPP usage by researchers at different types of institutions, the picture is a little different. Users at DOE labs account for 57 percent of the total MPP use, and university users account for 35 percent. Users at other labs and in industry each account for 4 percent of the MPP usage. These figures are based on FY00 use.



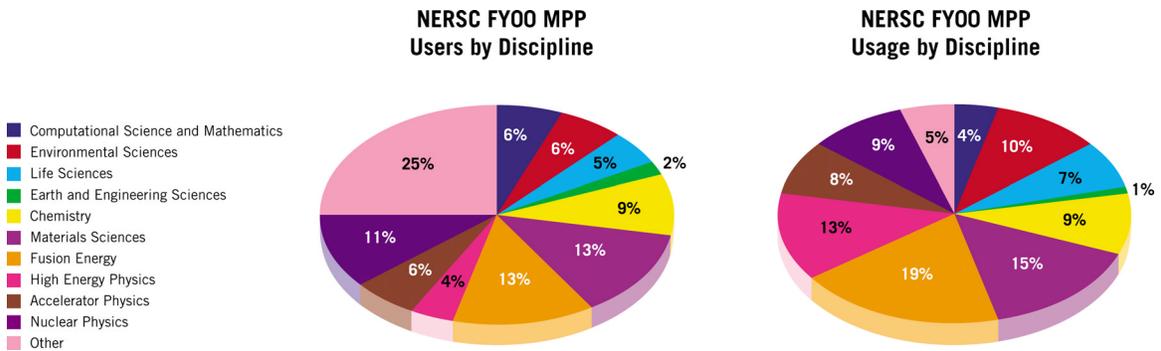
Use of NERSC's PVP computers is roughly equal between DOE laboratories and universities. Among users of the PVP systems, 46 percent are at DOE labs, 45 percent are at universities, 5 percent are from industry and 4 percent compute from other labs. As to PVP usage, 49 percent is by university researchers, 47 percent by scientists at DOE labs, and 2 percent each by industry and other labs.



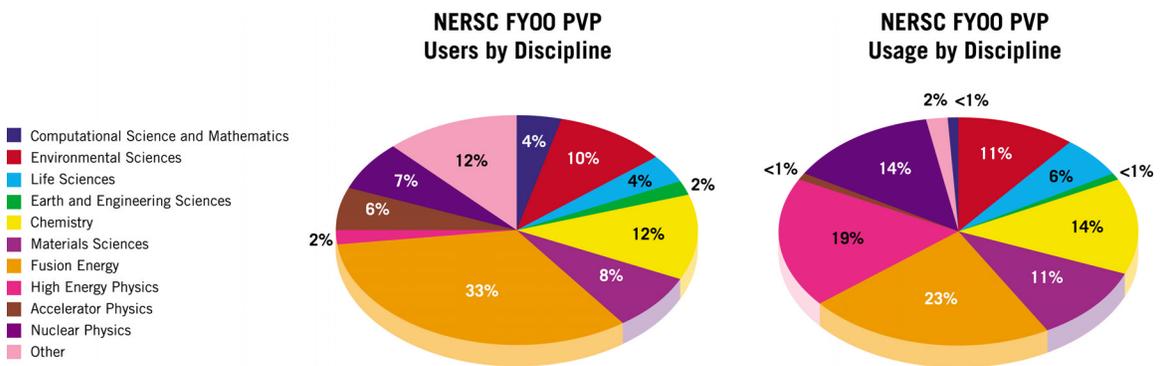
### By Scientific Discipline

NERSC supports research across the scientific spectrum of programs in DOE's Office of Science. The breakdown of the total number of users of NERSC's MPP systems by scientific disciplines shows Computational Science and Mathematics, 6 percent; Environmental Sciences, 6 percent; Life Sciences, 5 percent; Earth and Engineering Sciences, 2 percent; Chemistry, 9 percent; Materials Sciences, 13 percent; Fusion Energy, 13 percent; High Energy Physics, 4 percent; Accelerator Physics, 6 percent; Nuclear and Astrophysics, 11 percent; and Other, 25 percent. As to system usage by discipline, the breakdown is Computational Science and Mathematics, 4 percent; Environmental Sciences, 10 percent; Life Sciences, 7 percent; Earth and Engineering Sciences, 1 percent; Chemistry, 9 percent; Materials Sciences, 15 percent; Fusion

Energy, 19 percent; High Energy Physics, 15 percent; Accelerator Physics, 8 percent; Nuclear and Astrophysics, 9 percent; and Other, 5 percent.



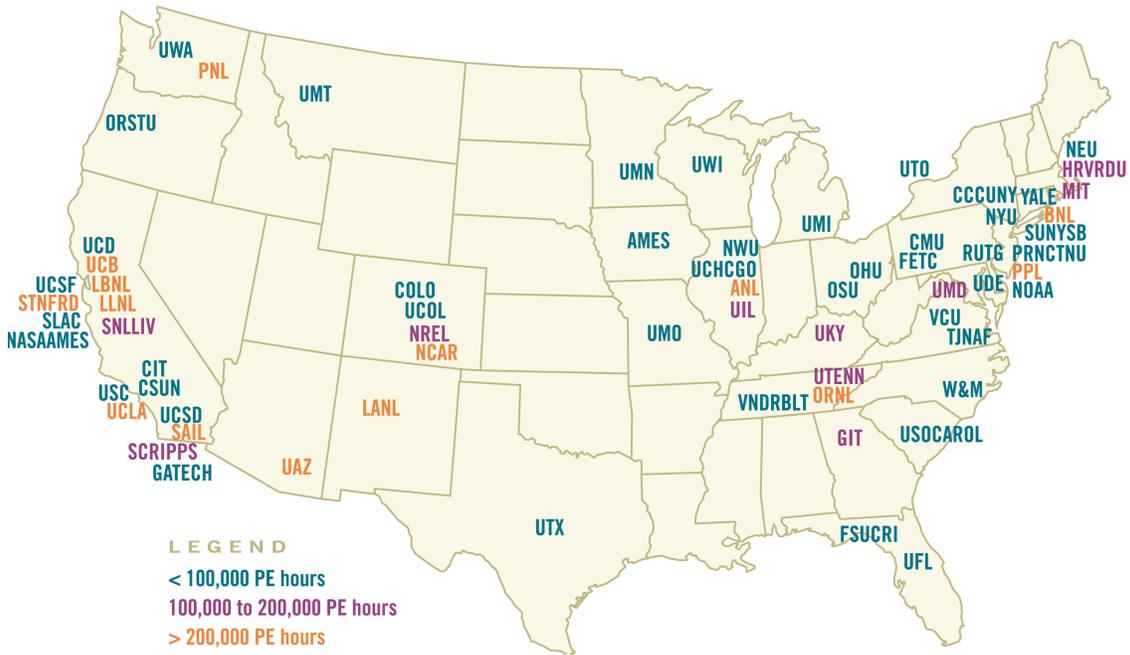
The distribution of PVP computer users by scientific discipline shows that the systems are most heavily used by fusion energy researchers, who account for 33 percent of the users and 23 percent of the total PVP usage. The remaining breakdown of the total number of PVP users by scientific disciplines shows Computational Science and Mathematics, 4 percent; Environmental Sciences, 10 percent; Life Sciences, 4 percent; Earth and Engineering Sciences, 2 percent; Chemistry, 12 percent; Materials Sciences, 8 percent; High Energy Physics, 2 percent; Accelerator Physics, 6 percent; Nuclear and Astrophysics, 7 percent; and Other, 12 percent. As to system usage by discipline, the breakdown is Computational Science and Mathematics, less than 1 percent; Environmental Sciences, 11 percent; Life Sciences, 6 percent; Earth and Engineering Sciences, less than 1 percent; Chemistry, 14 percent; Materials Sciences, 11 percent; High Energy Physics, 19 percent; Accelerator Physics, less than 1 percent; Nuclear and Astrophysics, 14 percent; and Other, 2 percent.



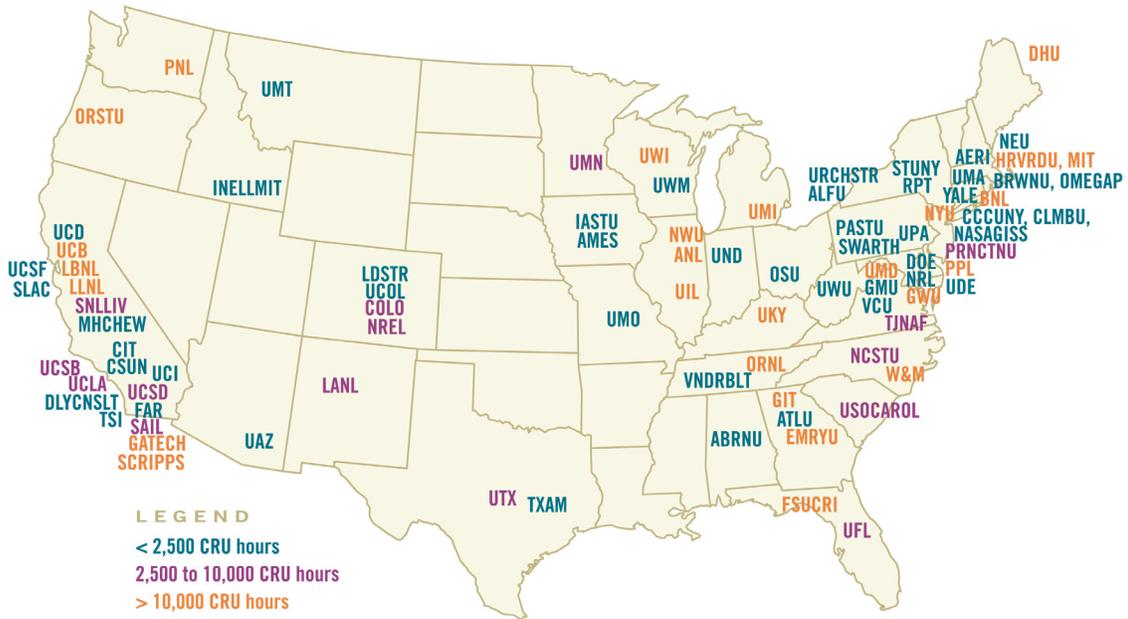
**By Geographical Distribution**

NERSC now serves users in 36 states. These maps show the distribution of NERSC clients, categorized by size of allocation, across the nation. This first map illustrates allocation of time on the massively parallel systems, while the second map shows institutions which have allocations on the PVP machines.

### MPP Allocations



### PVP Allocations



**A Leader by Any Measure**

We believe that when evaluated either as an organization supporting the Office of Science or compared with our counterparts supporting researchers funded by other federal and state agencies, NERSC will clearly stand out as a leading force in the world of high-performance computing. As one of the respondents to our FY2000 user survey put it, NERSC “Makes supercomputing easy.”

## RELIABLE AND TIMELY SERVICE

*Goal: Have all systems and support functions provide reliable and timely service to their clients.*

For the systems NERSC provides, service is assessed regarding availability, mean time between interruptions and mean time to repair computational and storage systems within six months of a system going into full service.

NERSC strives to provide reliable service to all of our clients. Our efforts address two general areas:

- How reliably our systems operate (i.e., availability to clients); and
- How responsive we are to clients when they have a problem.

To meet our goals, various groups within NERSC organization must work together to provide users with both the high-performance computing systems and the expert services for achieving research goals. To achieve this, NERSC takes a two-pronged approach. First, the NERSC staff is continually seeking out new techniques and technologies to anticipate and meet users' needs. Second, when a problem arises, we respond promptly to acknowledge, address and correct it.

One user who responded to our 2000 user survey summed up the results of our efforts by saying the NERSC "Provides excellent computing resources with high reliability and ease of use." Here are some examples of how we're achieving that goal.

**System Metrics For FY00  
FY 01 Goal (Measured in FY00)**

Systems	% Availability		Mean Time Between Interruptions (Hours)	Mean Time to Repair (Hours)
	Scheduled	Overall+		
Vector Systems	98.0 (99.36)	97 (99.11)	240 (423)	4.0 (3.8)
Storage Systems	98.0 (98.18)	97 (98.35)	120 (120)	4.0 (1.7)
Parallel Systems	96.0 (98.18)	95 (96.88)	96 (122)	4.0 (3.5)
Servers (fs/gw)	99.5 (99.52)	97 (99.39)	340 (2196)	8.0 (13.0)

NERSC strives to provide users with the maximum availability of our resources, not just in terms of scheduled availability, but in terms of overall availability. After all, if a system isn't available, it doesn't matter to the user whether it's a scheduled outage or unanticipated downtime. To ensure our systems are available, NERSC has set a goal of high availability on both a scheduled and overall basis, and as the chart above shows, we are exceeding our goal in both areas. Here are the system metrics definitions we use in setting our goals and evaluating our performance.

- Performance is shown with the goal listed first and the actual performance as measured in FY00 following in parentheses.
  - Scheduled availability is the percentage of time a system is available for users, accounting for any scheduled downtime for maintenance and upgrades.

$$\frac{\sum \text{scheduled hours} - \sum \text{outages during scheduled time}}{\sum \text{scheduled hours}}$$

- Overall availability is the percentage of time a system is running. In NERSC's 24 × 7 environment, 100 percent availability for FY00 (a leap year) would be 8,784 hours.

$$\frac{\text{available hours} - \sum \text{unscheduled outages and scheduled downtime}}{\text{available hours}}$$

- A service interruption is any event or failure (hardware, software, human, environment) that disrupts full service to the client base.
- Any partial degradation of committed services levels (e.g., dropping below the promised number of compute nodes on a system) is treated, for the sake of these goals, as a complete failure.
- Any shutdown that has less than 24 hours notice is treated as an unscheduled interruption.
- A service outage is the time from when computational processing halts to the restoration of computation (e.g., not when the system was booted, but rather when user jobs are recovered and restarted).
- If an outage occurs within two hours of the system that does not have checkpoint/restart being restored to service, it is treated as one continuous outage.
- If an outage occurs within two hours of the system being restored to service, it is treated as one continuous outage.

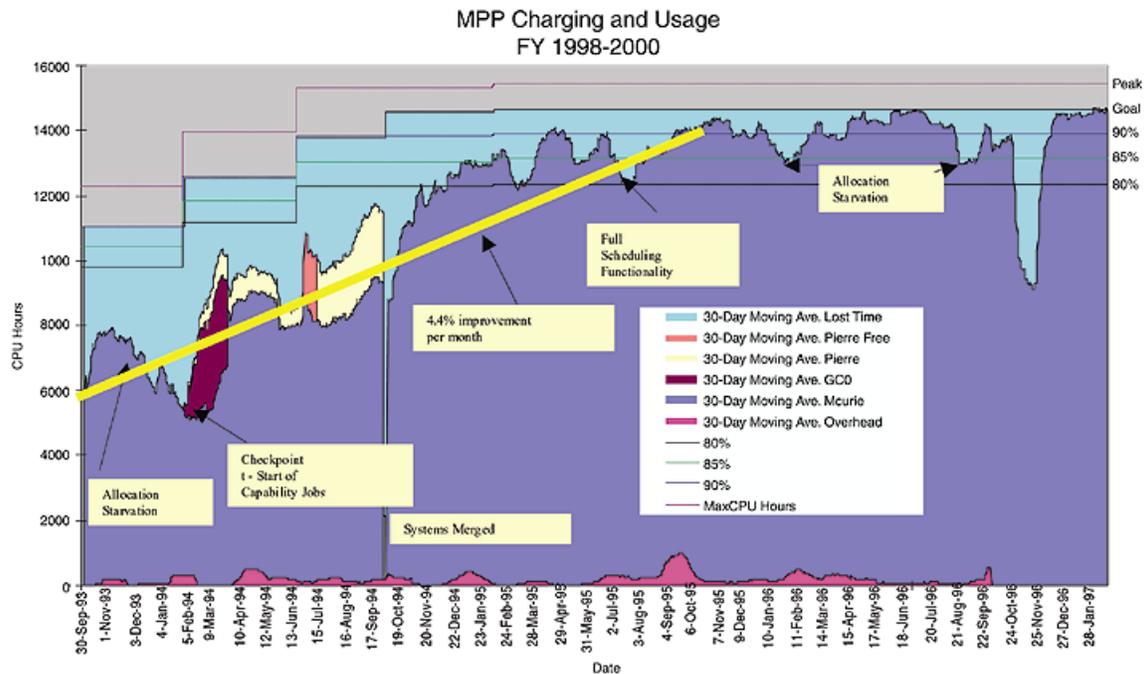
### **NERSC's Cray T3E Closing In on 95 Percent Utilization**

In early 2000, NERSC's Computational Systems Group took another step toward achieving a 95 percent utilization rate for the 696-processor Cray T3E. In April, a 30-day average rate of 94.56 was achieved, with just one day's rate of 92.9 percent keeping the group from reaching the 95 percent milestone. This achievement was due to the efforts of Tina Butler and Cray staffers Steve Luzmoor and Bryan Hardy, along with Mike Welcome's previous scheduler work. This required successful implementation of Cray Inc.'s "psched" scheduling daemon. With all the features of psched running, and with NQS and "prime job" control scripts written by Computational Systems Group staff, the T3E initially posted utilization figures of more than 93 percent, a level usually associated with capacity SMPs.

Efficiently scheduling a large MPP system is difficult. On the T3E, parallel applications are required to run on logically consecutive processors. Also, in the past, only one application could run on a range of processing elements in order to ensure synchronous scheduling. As applications entered and exited the system, the range of available processing elements would fragment, creating many small groups of unused processors. After a while, only small jobs could enter the

system. The psched load balancing feature automatically migrates running applications to collapse small holes of processors and thereby create large regions of available processors to run larger jobs. The psched gang scheduler will allow more than one application to run on a range of processors and will schedule them so that one application will run for awhile with complete control of the processors while the other is suspended. After a “time slice” is up, the applications switch roles and the suspended application gets to run. Another new feature of psched is the ability to designate a job with “prime” status. A prime job will preempt any other (non-prime) application and will be given preferred launching status to get in and running as soon as possible.

Here is an illustration showing the increasing utilization of NERSC’s Cray T3E.



### Utilization vs. Turnaround — A Tough Balancing Act

Achieving a high utilization rate for capacity jobs requires a tradeoff in other areas, such as the overall turnaround rate for all jobs in the system. Through the years, NERSC has worked to provide high utilization to meet the overall computing needs of its large and diverse user community and DOE programs. Higher utilization typically means higher allocations of resources. NERSC has tried to improve the balance between utilization and turnaround by creating special queues for specific jobs, but there is not one easy solution. In fact, just raising the issue of scheduling is a guaranteed means of generating a lively discussion at any meeting of computational scientists.

In trying to make improvements in scheduling, the approach must balance running interactive versus batch jobs, running multiple small sequential jobs as opposed to very large jobs, scheduling jobs that uses lots of processors versus jobs that don’t scale well. Adding to this mix are the limitations of the scheduling software. Each of NERSC’s computing systems uses a different

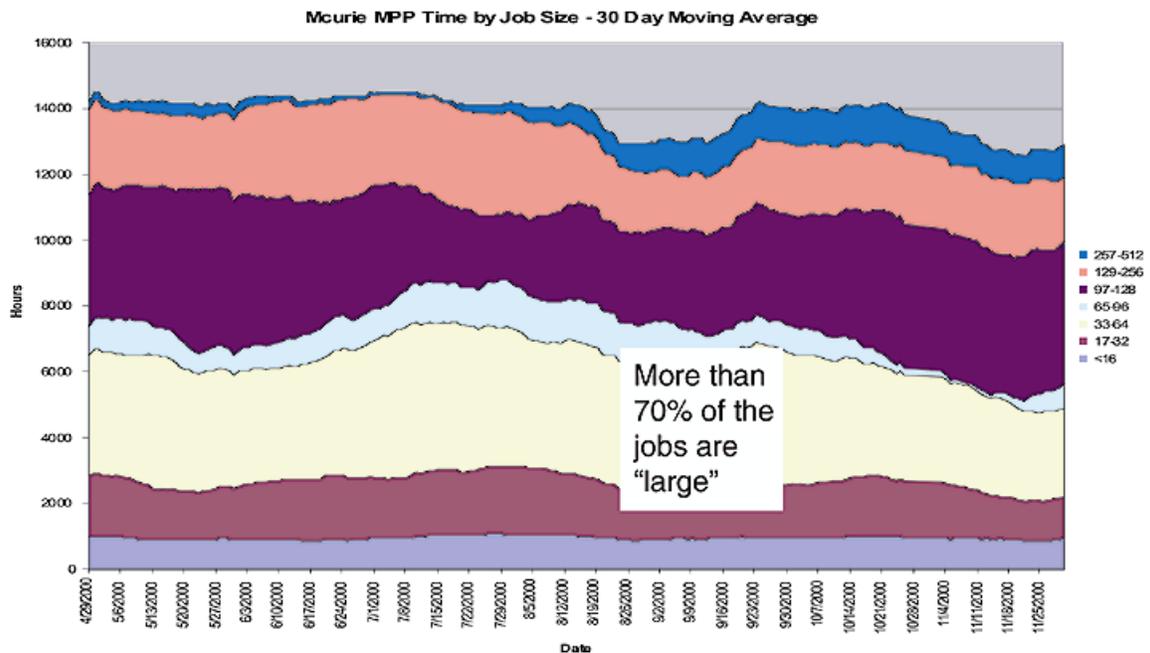
scheduling tool, so that a solution that works on one system may not be suitable for another machine.

As utilization of the T3E approaches 95 percent, the IBM SP Phase 1 averaged 85 percent utilization. This is due to differences in the scheduling systems of the two machines and the fact that it is easier to run large jobs on the T3E, and that the IBM automatically backfills unused processors with smaller jobs. The IBM utilization rate also has greater variation than the T3E.

NERSC realizes the importance of scheduling — it’s one of the subjects of most user comments on our surveys and in the 2000 survey, satisfaction with the turnaround rate on the T3E dropped. In 2000, due to comments from users, we created a long-running queue (12 hours maximum) for jobs using up to 256 PEs on the Cray T3E and this has decreased some of the demands from users. In 2001, NERSC plans to solicit ideas for scheduling improvements from our users and will report back on the results of this effort. Statistics for run and wait times for batch jobs on all NERSC systems are posted on the Web at <[http://hpcf.nersc.gov/running\\_jobs/qstat.html](http://hpcf.nersc.gov/running_jobs/qstat.html)>.

The job size chart above demonstrates that NERSC does run big jobs but also gives users a reasonable turnaround on smaller jobs most of the time. NERSC will continue to strive to achieve the most appropriate balance between the sometimes competing needs of our large user community.

### Job Distribution by Size on the Cray T3E



**Mass Storage Group Makes Transparent Move to Greater Reliability**

NERSC's Mass Storage Group made a systematic change in the production storage system, completely changing the underlying technology of the system — and NERSC users never saw a thing. The group consolidated the system from 22 standalone machines to 10 IBM SP nodes on two frames. To further improve reliability, the group also configured the storage environment to be completely redundant. As a result, any failure of a tape, disk or processor will not cause an outage of the entire system.

## **INNOVATIVE ASSISTANCE**

*Goal: NERSC aims to provide its clients with new ideas, new techniques and new solutions to their scientific computing issues.*

NERSC is committed to helping its clients achieve better performance and results from their computational science efforts. Merely providing the computers, data storage and software is not enough to do this, so NERSC has initiated a range of activities to provide innovative assistance. From individual outreach to general tutorials, these programs demonstrate NERSC's commitment to being an integral part of research conducted by Department of Energy Office of Science programs. This assistance can take on many forms and here are a few examples.

### **NERSC Introduces New IBM SP Via Videoconference Training**

In March 2000, member of the NERSC User Services Group held a training videoconference to acquaint users with NERSC's new IBM SP System, [gseaborg.nersc.gov](http://gseaborg.nersc.gov), which was slated to go on line in April. The training was distributed to other national labs via ESnet's videoconferencing service. This presentation described the system and provided pointers to more detailed information for users who were new to SP systems. Also, consulting staff were on line to answer specific questions from attendees.

NERSC is continuing to investigate other on-line training technologies, including point-to-point netcasting.

### **BIO — a Parallel Application for Deciphering the Human Genome**

As the Human Genome Project enters its next phase — decoding the massive amount of gene sequences already produced — NERSC staff are helping to create an assembly code that can keep pace with the output of DNA data from the newest generation of sequencing machines. Currently, there is no publicly available code that can handle the computational and memory requirements for performing large-scale genome assembly.

That's where NERSC's Leonid Olikier and Sherry Li come in. Together, they are working with Dan Rokhsar's group at DOE's Joint Genome Institute in Walnut Creek, Calif., to develop BIO, a code able to run on parallel high-performance computers. The combination of such a code and the latest supercomputers, such as the newly installed NERSC-3 system, will allow scientists to complete their analyses in a reasonable amount of time. Once the DNA assembly is complete, the next step will involve identifying individual genes and discovering their biological functionality. Solving these multidisciplinary problems will require radically new computational methodologies.

Developing BIO poses several challenges. First, it is being developed at the same time as the serial computing code, rather than adapting an existing serial code for a parallel platform. Other challenges involve designing parallel algorithms that can efficiently handle the irregularities of the data structures, as well as in the actual code implementation. In particular, the core

combinatorial algorithms are very hard to implement scalably on a large-scale parallel machines. The software should be able to handle mammalian-level assemblies, comprised of billions of DNA bases, by the end of the year 2001. The human genome sequence consists of about 3 gigabases.

The development of such bioinformatics tools as BIO reflects a new twist in the field of high-performance computing. Until now, most of the development of both systems and applications was driven by the fields of physics and engineering. With the explosion of data and computational requirements from such efforts as the Human Genome Project, bioinformatics is becoming a key application area.

### **Helping Researchers BLAST Through Mountains of Genome Data**

Computational biology, still an emerging field, is already being swamped by its own success. On the one hand, the Human Genome Project and similar efforts to identify DNA of other organisms have produced vast amounts of publicly available data. Conversely, new technologies for sequencing genomic information are adding to existing mountains of data at an unprecedented rate. The problem now facing genomic researchers is how to get the science out of this data.

Jonathan Carter of NERSC's User Services Group, working in conjunction with DOE's Joint Genome Institute, is parallelizing BLAST, the Basic Local Alignment Search Tool as implemented by the National Center for Biotechnology Information. Adapting BLAST to run on parallel platforms will allow scientists to more quickly search the existing data for useful information.

Research groups around the world are obtaining the genetic structure of different organisms. This information is being obtained and distributed in quantities ranging from whole genomes to individual genes to proteins to fragments containing known, interesting biological characteristics. This data is being posted in various genome databases around the world and can be downloaded for analysis. After a genome has been assembled, scientists want to know where the individual genes are and what they do. By comparing new sequences against other databases, researchers can identify new genes, or help identify posted segments from unidentified organisms. This can be both time- and computing-intensive when an organism has 30 million base pairs of genetic information and the researcher wants to find all similarities between this data and that of another organism.

The BLAST application compares a new sequence with stored databases and takes into account gaps in the data and not-quite-perfect matches. The results are then listed by scoring similarities between the new sequence and the database information. In performing this comparison, BLAST takes into account gaps in the genetic information or the insertion of new pieces of genetic material.

Jonathan is currently working to adapt the code so it can run on NERSC's 3.75 teraflops IBM SP system, which will provide scientists with a powerful computing resource for decoding the basic building blocks of life.

### **That Majdi Touch**

Although NERSC and other HPC centers are often characterized by the lineup of hardware they provide to users, NERSC takes particular pride in the human side of the services we offer. Our staff routinely communicate with users on the phone, in person, via email and Web interfaces. Usually the question is answered or the problem resolved and both the user and consultant move on to the next job. Occasionally, however, a user takes the time to point out special or innovative service. In December, Rui Li, a staff scientist at Jefferson Lab, wrote to praise the assistance provided by Majdi Baddourah of the User Services Group. Here is an excerpt from that message:

“I developed a code studying the effect coherent synchrotron radiation on beam dynamics, this is so far the only self-consistent simulation of this problem which is crucial to many of the next generation machine designs. Over the past 3-4 years, Majdi has helped me in optimization, debugging, converting to MPI, transfer to SP machine. Because of his help, the program is very well received in the accelerator physics community, has attracted international collaborations, and received funding from DOE’s grand-challenge program. Without Majdi’s help, all these achievements wouldn’t be possible.”

“I’m grateful that you have such a capable and helpful staff in the group that the users can rely on with confidence.”

## TIMELY AND ACCURATE INFORMATION

*Goal: Provide timely and accurate information and notification of system changes to the client community so they can most effectively use the NERSC systems.*

The NERSC staff strives to provide our clients with timely and accurate information which may affect those clients' research efforts. Not only do we give adequate notice of changes and outages, we also try whenever possible to provide an explanation of the reasons behind the change and the expected impact on clients. As an example of timeliness:

- Planned system changes were announced at least seven days in advance (except in one instance); all planned system outages announced at least 24 hours in advance.
- All system changes and planned outages were announced in advance on the NERSC "What's New" Web page. Some major changes were also announced by email to PIs. By combining web postings and email announcements, we are able to more quickly inform clients of changes, which in turn allows us to implement changes more quickly.

In 2000, NERSC User Services Group instituted a notification program under which users could sign up to receive specialized email messages for a variety of topics, allowing them to tailor the information they receive. Also, as a result of findings from the 1999 survey, several other changes were made on NERSC's HPCF website. The changes made it easier to find information on the web concerning running batch and interactive jobs and a chart showing automated machine up/down information was placed on the bottom of HPCF Home Page. Also, a NERSC Glossary and Acronym List was added to the site.

These changes, combined with the revamped NERSC Web pages which debuted in 1999, continue NERSC's goal of improving the flow of timely and accurate information.

As one respondent to the 2000 survey indicated, "The announcement managing and web-support is very professional."

### Getting Ready to Run on the SP

In addition to supporting NERSC users on the Cray systems, NERSC staff focused strongly on helping our user community get ready to run their codes on the IBM SP system, the first phase of which arrived in summer of 2000.

In April 2000, NERSC sponsored "A Workshop on the IBM SP System" presented by IBM experts. John Levesque's staff from IBM's Advanced Computing Technology Center (ACTC) presented three mornings of lectures and three afternoons of workshops on porting and optimizing code for the SP. Lectures covered cover all aspects of SP system hardware and software, including Power3 processor architecture; development, compilation, debugging and tuning tools; and techniques for achieving high performance in serial and parallel execution. The instructors for this workshop were Charles Grassl, Bob Walkup, and David Klepacki.

At the June NERSC Users Group meeting held at Oak Ridge National Laboratory, the IBM SP was featured in several staff presentations. These talks included:

- The IBM SP: Evolution from Phase I to Phase II: Described the IBM SP system in its current state and what it will be after enhancement. System configuration and expected performance were outlined.
- Profiling Tools on the NERSC Crays and SP: An overview of performance analysis and profiling tools available on the NERSC Cray and IBM SP platforms. Modifications (e.g., compile, link and runtime options) required to access the tools were outlined and sample results presented. Particular emphasis was placed on the current status of profiling tools on the NERSC IBM SP.
- Tricks and Tips for Porting between the T3E and the SP: NERSC consultants have come up with a number of tricks and tips for porting programs written for the T3E to the SP and this talk covered some of the most useful tips and techniques.

### **Getting the Word Out About NERSC Systems**

NERSC uses several methods to inform users about changes and updates in its systems. These methods include email notification, Messages of the Day (MOTD) posted on the NERSC High Performance Computing Facility website, and announcements pertaining to specific systems.

Systems announcements: For the first half of 2000, changes and updates in specific systems were announced by posting them on the system-relevant web pages and on the Web. In August, NERSC began making such announcements via email. During the year, a total of 233 such announcements were made. The number of postings for each area were: 97 notices on the status of machines, 24 notices for the IBM SP, 28 notices for the Cray T3E, 22 announcements for the PVP (J90 and SV1) systems, 18 notices for file storage, 17 server updates, 16 general announcements, two notices for the Cray Programming Environment, three notices for the Programming Library, and six announcements for Software Applications and Tools.

## **NEW TECHNOLOGIES, EQUIPMENT, SOFTWARE, AND METHODS**

*Goal: Ensure that future high-performance technologies are available to Office of Science computational scientists.*

From the outset, NERSC has provided its users with some of the most up-to-date computing resources available. In mid-2001, NERSC will offer its users the world's most powerful unclassified computer when the new 2,528-processor IBM SP goes on line. As new machines are installed, systems experts carefully analyze performance to ensure that the equipment meets the high-performance needs of NERSC clients, and to continually work with vendors to increase reliability and performance.

The process of meeting users' needs begins anew with each procurement cycle, with each step clearly spelled out. The NERSC Users Group, or NUG, outlines the basic requirements in a document known as the "Green Book." After reviewing the Green Book laying out requirements for a new system, NERSC staff representing various groups then produce a request for proposals and critically subject each vendor's proposal to careful review. Before NERSC adopts a new technology, the equipment must be fully evaluated and tested. (See accompanying section on Wise Technology Integration.)

One way of doing this evaluation and test is by using NERSC's infrastructure as a testbed. Here are a couple of examples.

### **NERSC, IBM Make Smooth Phase Transition with New SP**

When NERSC announced that it had selected an IBM SP system for its next HPC resource, the contract called for the machine to be installed in two phases. As a result, a 608-processor "Phase 1" SP was installed at Berkeley Lab during summer 2000. As part of the acceptance testing, a select group of NERSC users were given access to the system. This allowed NERSC staff to assess the workaday performance of the system, while allowing some users to begin moving their jobs to the new system. During this time, NERSC was able to integrate the SP platform into the center and to provide the necessary support expertise, so full support could be provided to users once the larger "Phase 2" system is installed in 2001. This two-step implementation allowed IBM more lead time for delivering a system with higher performance capabilities while also allowing NERSC users to better prepare their applications to take greater advantage of the system. As part of the procurement, IBM and NERSC are also working to develop more real-world benchmarks for evaluating system performance.

### **NERSC Continues to Build Expertise in Cluster Computing**

Another system already in place but being constantly improved upon is the 264-processor PDSF (Parallel Distributed Systems Facility) cluster dedicated to high energy and nuclear physics. Originally part of the now-canceled Superconducting Supercollider, PDSF arrived at NERSC as a collection of networked workstations. With judicious upgrading of hardware and software, the PDSF has been quietly going about its task of supporting HENP research for the past three years,

during which many other organizations have announced plans to purchase or develop similar cluster computers for specific scientific disciplines.

Based on our success with PDSF and experience with two other clusters used for research, Berkeley Lab has purchased a 160-processor cluster to assess whether such a system can meet the day-to-day production demands of a scientific computing center. To date, most large scientific commodity clusters are used more for specific research applications than general-purpose production resources for computational science.

Although clusters are often heralded as the next wave in high-performance computing, they are just now coming into their own. The goal for NERSC, which will operate the new Berkeley Lab cluster for a specified period, is to see how well a commercially built cluster can perform in a demanding scientific environment. The cluster will consist of 160 Intel Pentium III processors in 80 nodes, with one gigabyte of main memory in each node. Each node will be connected by a high-performance interconnect from Myrinet. The cluster will have a total of 500 gigabytes of shared disk space.

The primary goal in purchasing the cluster system is to provide a system testbed for NERSC and other staff to explore the applicability of such a system to support a highly parallel, numerically intensive workload. NERSC is interested in evaluating cluster architectures as an option for procurement of large production computing systems in the future. If successful, the cluster architecture could be a candidate for procurement of future NERSC production systems.

### **Probe HPSS Project Provides a Valuable Technology Testbed**

Last year's "How Are We Doing" report described Probe, a DOE-funded collaboration between NERSC and Oak Ridge National Laboratory to develop a testbed for storage-intensive applications. Unlike some testbeds which consist of cast-off equipment, the Mass Storage Group's Probe testbed is a state-of-the-art environment for evaluating new technologies related to HPSS (the High Performance Storage System). Because of its advanced technology, Probe is an attractive facility for vendors to test new technologies. For example, in fall 2000, IBM beta-tested its new LTO tape technology and library. The new technology offers a tape capacity of 100 gigabytes of compressed data, or 300 gigabytes uncompressed. The new technology was tested on Probe and the ability to mount it in HPSS was demonstrated in October. IBM now plans to roll out the technology commercially in late 2001.

The Mass Storage Group also installed and configured Metastor's E4400 full-fibre disk in NERSC's production systems in June. The group replaced all of the HIPPI-attached network storage devices with this new capability, which is resulting of disk speeds of 80-90 megabytes per second.

## WISE TECHNOLOGY INTEGRATION

*Goal: Ensure all new technology and changes improve (or at least do not diminish) service to our clients.*

Any given IT trade show these days features dazzling displays by companies that weren't in existence three years ago and may not be around for next year's expo. It's not enough anymore to just buy the latest technology. Centers like NERSC must ensure that the technology won't disrupt the systems and services our users depend on. To ensure that the systems it provides to clients are reliable, NERSC thoroughly evaluates the technology and tests systems off-line before deciding whether to implement the system for production or file the results away under "lessons learned." A common aspect of the later stages of such evaluations is inviting specific NERSC users with demanding applications to help put systems through their paces to see how equipment stands up to real-world demands. Feedback from users then helps the staff further improve the systems and services. Here are some examples of this wise technology integration as practiced at NERSC during the past year.

The most far-reaching integration of new technology at NERSC during 2000 was the installation of Phase 1 of the new IBM SP system. This was the first non-Cray machine ever to be deployed by NERSC as a major computational resource and required both the staff and the user community to adapt to a new platform. By carefully and precisely detailing the system requirements during the procurement and then working closely with IBM once the new machine arrived, NERSC was able to provide users with a reliable next-generation system. In fact, the implementation went so smoothly, NERSC was able to allocate an additional 1 million hours to users of the IBM SP. And when the Cray T3E was taken off line for the move to the new Oakland Scientific Facility, the new IBM remained in service, providing the user community with uninterrupted service.

As one of the respondents to our survey noted in fall 2000, "NERSC has been the most stable supercomputer center in the country particularly with the migration from the T3E to the IBM SP."

### **Effective System Performance (ESP): A Real-World Benchmark**

HPC vendors and centers alike tout their systems according to theoretical peak performance figures, but these are not true indicators of what a high-performance computing system can actually provide. To provide more realistic performance evaluations of HPC systems in the workaday world of scientific computing, NERSC staff members Adrian Wong, Lenny Oliker, Bill Kramer, David Bailey and others developed a new benchmark test called Effective System Performance, or ESP.

In short, the test measures both how much and how often the system can do scientific work. ESP, which is easier to use than other benchmarks, gives HPC centers a tool for bridging the gap between new technology introductions and wise technology integration. Two other centers, the Secure Computing Facility at Lawrence Livermore National Laboratory and the National Center

for Atmospheric Research in Colorado, have already adopted ESP as a benchmarking tool, and NERSC has also been working with interested vendors regarding the application of ESP.

The traditional benchmarking approach includes simple CPU performance tests, a small number of pseudo-applications, an I/O test and, sometimes, a structured throughput workload test. However, this approach does not measure such real-world processes as integrated system function, jobs with varying degrees of CPU, memory and I/O requirements, random mixes of jobs, system administration/resource management, system behavior, usability, or slowdown vs. utilization.

ESP is valuable in that it measures system effectiveness, which is a significant factor in determining the production capability of a system. This benchmark, which has been run on NERSC's Cray T3E and IBM SP, provides important information on how a system can be made more effective through scheduling. Improving system effectiveness to increase the number of jobs being run is comparable to adding more processors, but can be done without the additional hardware costs.

ESP is designed to evaluate systems for overall effectiveness, independent of processor performance. The ESP test suite simulates "a day in the life of an MPP" by measuring total system utilization. Results take into account both hardware (PE, memory, disk) and system software performance. Developed by NERSC as part of a major system procurement process, ESP is designed to predict the effectiveness of a system before purchase, as well as to evaluate system changes before implementation.

The goals of the ESP benchmark include:

- determine how well an existing system supports a particular scientific workload
- assess systems for that workload before purchase
- provide quantitative effectiveness information regarding system enhancements
- compare different systems on a single workload or discipline
- compare system-level performance on workloads derived from different disciplines
- compare different systems for different workloads.

### **HSI Interface for HPSS Is Easier to Use, Provides More Utility — and Runs Faster**

In developing the HSI interface for NERSC's High Performance Storage System (HPSS), the Mass Storage Group sank a three-pointer for an all-around win in wise technology integration. The HSI interface client was developed using the Probe HPSS testbed (for more information, see the section on Technology Transfer). This allowed the new technology to be thoroughly tested in a full-blown, state-of-the-art HPSS environment for an accurate demonstration of how well the client would perform in the production HPSS system. The HSI was developed to provide users of the former CFS storage system with an easier-to-use interface. Not only is HSI more user friendly, it's also better in that it allows the user know the status of a data transfer, such as whether a transfer is successful or not. The new interface also offers a 10-fold increase in data transfer rates, up from 3 megabytes per second to 30 megabytes per second.

### **The Wisdom of Collaborating with Vendors**

As also mentioned in the New Technologies section, the Probe testbed served as a beta test site for the new IBM LTO technology tape drives and library. This allowed IBM to gain a thorough evaluation of the technology — again in a fully operational environment — before rolling it out for wider use.

Finally, one way to make the transition to new technologies work smoothly is to help provide input from the ground up as new systems are developed, and NERSC is well represented in this regard. Bill Kramer is a member of both the IBM and Cray advisory boards, while David Bailey is a member of the Cray Inc. advisory board for the development of the Cray SV2 computer.

## PROGRESS MEASUREMENT

*Goal: As a national facility serving thousands of researchers, NERSC has a responsibility to our clients to measure and report on how we're doing in terms of providing service, support and facilities.*

As the nation's first unclassified supercomputing center and the DOE's flagship high-performance computing center for unclassified research, NERSC has long been held to high standards. Since moving to Berkeley in 1996, NERSC has begun assessing our own performance against a set of goals — and reporting back to users and stakeholders. This report, now in its fourth edition, is one of several documentation efforts on how we're doing. We also report — and get feedback — regularly at biannual NUG meetings and post statistics covering a wide range of our systems and services on the Web.

### Surveying Our Users

For the third consecutive year, NERSC surveyed its user community in 2000 to find out what works well and what could be improved. As with the previous surveys, the results are carefully considered and, when appropriate, used as the basis for changing services to users. For example, as a result of the FY99 survey, NERSC staff made the following changes:

- Created a long-running queue (12 hours maximum) for jobs using up to 256 PEs on the Cray T3E. Last year, seven users asked for longer T3E queues; this year only one did.
- Opened a Cray SV1, seymour, for interactive use — a change well appreciated by users.
- Created new email lists to keep users better informed of NERSC announcements and changes. (This change wasn't reflected in this year's survey results.)
- Enhanced the HPCF website; overall satisfaction with the website was higher this year. Made it easier to find information on the web concerning running batch and interactive jobs.
- Added a NERSC Glossary and Acronym List.
- Placed automated machine up/down information on the bottom of HPCF Home Page.

### NERSC FY 2000 User Survey Results

In FY 2000, 134 users responded to our survey. The respondents represent all 5 DOE Science Offices and a variety of home institutions. Here is a snapshot of the results:

On a 7-point scale, with 7 corresponding to Very Satisfied and 1 to Very Dissatisfied, the average scores ranged from a high of 6.7 for our training classes and the PVP Fortran compilers to a low of 4.3 for PVP and T3E batch job wait time.

Other areas with very high user satisfaction are consulting advice and SP availability (uptime).

The areas of most importance to users are the available computing hardware (the amount of cycles), the overall running of the center and its connectivity to the network.

This year, the largest increases in user satisfaction came from the PVP cluster. Following the conversion of Seymour last year to an interactive machine, user satisfaction for the ability to run interactively on the PVP increased by almost one point. Five other PVP ratings increased by 0.6 to 0.8 points. See the hardware and software sections. Other areas showing a significant increase in satisfaction are HPSS performance and response time, hardware management and configuration, the HPCF website, and the T3E Fortran compilers. Only two scores were significantly lower this year than last: The first change was in consulting, which received the highest rating last year and remained very high this year. The second is for the T3E batch wait time, which in part reflects the pent-up demand for more computing resources.

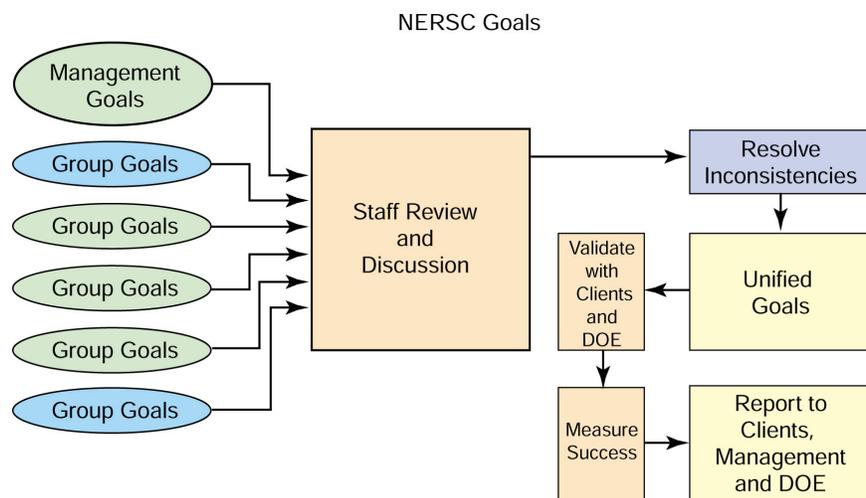
When asked what NERSC does well, 34 respondents focused on NERSC's excellent support staff and 29 pointed to our stable and well managed production environment. Other areas singled out include well-done documentation, good software and tools, a very useful storage environment, and well managed migrations and upgrades that "make supercomputing easy."

When asked what NERSC should do differently, the most common responses were to provide more resources, especially more cycles and nodes.

Of the 47 users who compared NERSC to other centers, 53 percent said NERSC is the best or better than other centers.

### Self-Assessing Our Systems and Services

Just as our progress reporting is an interactive and public process, so too is the process of setting and validating our goals. Ideas for the goals originate both with NERSC management and staff. The ideas are then discussed, combined and refined at all-hands meetings. Once agreed upon by our organization, the goals are then presented to the NERSC Users Group for discussion and validation. The resulting version is then presented to stakeholders in DOE for their consideration. Only then are the goals implemented for use in assessing our systems and services. Here's a graphic representation of how that process works.



## **New Staff Goals for 2001**

In the year 2000, NERSC staff reviewed the goals originally adopted as the organization was getting established at Berkeley Lab. After extensive discussion and some revision, a new set of goals was adopted for use beginning in year 2001. Many of the goals are only slightly changed, some have been combined and others added. These goals, outlined below, will serve as the framework for next year's "How Are We Doing?" report.

### **Systems and Service Goals from the Client's Viewpoint**

- 1. Reliable and Timely Service.** For the systems NERSC provides, service will be assessed regarding availability, mean time between interruptions and mean time to repair computational and storage systems within six months of a system going into full service.
- 2. Client Support Goals.** The end measure of a site is how much productive scientific work users accomplish. Sites must assist users in being as productive as possible by providing systems, tools, information, consulting services and training. The objective is to understand codes and how they are used, and target bottlenecks for elimination or minimization.
- 3. Never Be a Bottleneck to Moving New Technology into Service.** NERSC is a primary vehicle for achieving the SC goal of making leading-edge technology available to its scientists. To do this, NERSC continually evaluates, tests, integrates and supports early systems and software. Therefore, NERSC must help ensure future high-performance technologies are available to Office of Science computational scientists in a timely way.
- 4. Ensure All New Technology and Changes Improve (or at Least Do Not Diminish) Service to Our Clients.** In striving to provide users with the latest systems for computational sciences, NERSC has the responsibility to ensure system changes have a maximum benefit and minimal detrimental impact on the clients' ability to do work.
- 5. Develop Innovative Approaches to Help the Client Community Effectively Use NERSC Systems.** NERSC must assist our clients in being as productive as possible by providing systems, enhancements, tools, information, training, consulting and other assistance. In addition to the traditional approaches that are effective, NERSC will constantly try new approaches to help make our clients effective in an ever-more-changing environment. NERSC will help design strategies and integrate and develop technology to enable our clients to improve their use of our systems and to more effectively accomplish their science.
- 6. Develop and Implement Ways to Transfer Research Products and Knowledge into Production Systems at NERSC and Elsewhere. (New for 2001)** NERSC is uniquely placed to establish methods and procedures that enable research products and knowledge, particularly those developed at LBNL/UC, to smoothly flow into production.

**7. Improve Methods of Managing Systems Within NERSC and LBNL and be a Leader in Large-Scale Systems Management and Services.** As the Department of Energy's largest unclassified scientific computing facility, NERSC continually provides leadership and helps shape the field of high performance computing. As HPC technology evolves at an increasing rate, it is crucial that NERSC and LBNL remain at the forefront of getting the most out of these systems.

**8. Export Knowledge, Experience and Technology Developed at NERSC, Particularly to and Within NERSC Client Sites.** In order for NERSC to be a leader in large-scale computing, NERSC must export experience, knowledge, and technology. Transfer must be made to other client sites, supercomputer sites, and industry.

**9. NERSC Will Be Able to Thrive and Improve in an Environment Where Change Is the Norm.** (New for 2001) High-performance organizations that deal with advanced technology must be able to adapt and embrace change as a way of life. HPC centers that are not growing and changing are dying (or have died). Providing reliable cycles is not enough to serve the NERSC users in a time of constant change. Research is needed to ensure that tomorrow's systems are accessible and productive to our users.

**10. Improve the Effectiveness of NERSC Staff by Improving Infrastructure, Caring for Staff, Encouraging Professionalism and Professional Improvement.** Every employee has a stake in the success of NERSC and management encourages staff to contribute their ideas for helping the organization succeed. To help facilitate the professional exchange of ideas and information, NERSC has adopted a series of guidelines and information. They are posted at <<http://www.nersc.gov/staff/#nersc>>.

## HIGH-PERFORMANCE COMPUTING CENTER LEADERSHIP

*Goal: Improve methods of managing systems within NERSC and be the leader in large-scale computing center management.*

As the Department of Energy's flagship unclassified scientific computing facility, NERSC has continually provided leadership and helped shape the field of high performance computing since the center was established in 1974. This pattern continued in 2000, as the NERSC staff's experience and expertise was demonstrated in any number of ways, from creating a state-of-the-art computing facility in Oakland to shaping the software infrastructure for international research. Here are some examples of NERSC's expertise and knowledge-based leadership in enhancing large-scale computing in 2000:

### **New Oakland Scientific Facility's Computer Room Designed for Flexibility, Expandability**

When it became apparent in 1999 that the Laboratory would need a new facility to house the NERSC systems, in addition to other Lab programs, a set of technical requirements were laid out, forming a blueprint for future adaptability and expansion. In constructing the new machine room, NERSC implemented some of its own innovative ideas, as well as drawing on best practices from other centers. The result is a facility able to accommodate any high-performance computing system expected to be developed in the next 20 years.

Specifications for the new machine room included a minimum of 20,000 square feet of floor space, with an option for an additional 20,000 square feet of contiguous space. The room was to be 19 feet from true floor to true ceiling, with a three-foot raised floor capable of supporting 250 pounds per square foot. The building needed to meet current California seismic requirements, have 5 megawatts of electrical power and be connected as a node on ESnet (DOE's Energy Sciences Network), as well as be convenient to public transit. Finally, the building had to be ready for occupancy by fall 2000 and be secured through a 10-year lease with options to renew. The Oakland Scientific Facility (OSF), located in downtown Oakland, meets or exceeds all of these requirements.

Although the first-floor machine room in the OSF has all the appearances of a state-of-the-art facility, many of the features which make it so are hidden from view. The three-foot raised floor, twice the height of the old NERSC machine room at Berkeley Lab, is required to accommodate the cabling requirements of the newest high-performance computing systems. Unlike the specially built supercomputers of the 20<sup>th</sup> century, newer systems are larger and require more interconnects. To facilitate such connections, the raised floor at OSF (which meets current seismic standards) was built using only vertical and horizontal supports, instead of the more common diagonal supports. This approach makes it easier to route cables, conduit and pipes, and also results in a floor that is seismically stronger. In addition to being simpler and less expensive, this technique allows 50 percent fewer direct seismic ties for the computer cabinets.

The new center also uses a simplified and more flexible power distribution system, utilizing power distribution units (PDUs) on the machine room floor. The PDUs incorporate circuit breakers, transient noise filters and voltage monitors and eliminate the need for fixed conduit connections to each machine. This allows computers to be moved and quickly reconnected to the power supply without having to rewire the machine room. In short, it makes power connections faster, less expensive and more flexible.

The entire machine room is overseen from a centrally located control room. In addition to providing an all-encompassing view of the facility, the machine room also serves as the initial monitoring center for ESnet.

After more than a year of planning, months of construction, and an intense week of moving and installation, NERSC's Cray, HPSS, PDSF and auxiliary systems were quickly and efficiently moved over a long October weekend from Lawrence Berkeley National Laboratory to Berkeley Lab's new Oakland Scientific Facility. The move was completed ahead of schedule and the systems were back on line soon than anticipated. The move, like the 1996 move of NERSC from Lawrence Livermore to Berkeley Lab, was completed without a break in service to users as NERSC's IBM SP Phase I system remained in service during the transition.

### **David Quarrie Takes on Job of Chief Architect for ATLAS Computing**

David Quarrie, leader of NERSC's High Energy and Nuclear Physics Computing Group, has accepted a two-year assignment as chief architect for the ATLAS high-energy physics experiment. ATLAS is an international research program to be carried out at the Large Hadron Collider at CERN in Switzerland beginning in 2005. David's new position will occupy about 50 percent of his time, and he will also continue working on software efforts related to the BaBar project at the Stanford Linear Accelerator Center.

The job of the chief architect is to establish a coherent vision for the software, David said. On a practical level, the task is to produce and implement the framework or environment in which scientists will write the physics algorithms to do the physics they need to do as part of the ATLAS experiment. The architect is responsible for all the off-line software for reconstruction, simulation, physics analysis and the final real-time filter for on-line data generation and collection. The job requires producing results and meeting a series of milestones every three to four months over the next several years. David's team met its first milestone by producing a prototype of the framework.

David accepted the position after serving as a member of the ATLAS Architecture Task Force. This will be the third attempt to produce the software architecture and David is hoping to make it a success. He says the challenges are more sociological than technical. ATLAS involves about 1,800 researchers from dozens of countries.

### **NERSC's Bill Kramer Leads Construction of SC2000's Massive Network**

In the months leading up to the SC2000 Conference in Dallas last November, Bill Kramer was wearing another hat — a 10-gallon hardhat — in addition to his role as head of NERSC's High Performance Computing Department. As Vice Conference Chair for Information Architecture, Bill led the team charged with designing, building and operating SCinet, the SC conference network. When they were done, they had created one of the nation's largest communications networks consisting of three OC-48 lines and three OC-12 lines. Bottom line: SCinet offered a combined peak speed of 9.4 gigabits per second.

As Bill explained in a news release released during the conference, "Because the SC conference is centered around the latest achievements in high-performance computing, it's only fitting that we build one of the world's biggest network connections to support it. At last year's conference in Portland, SCinet provided more connectivity than all the combined networking resource in the states of Oregon and Washington — and this year's version of SCinet is even bigger and faster." Conference attendees also agreed that the 2000 version of SCinet was the most robust ever, with many fewer lost connections and outages.

Qwest Communications, Inc. provided the main networking connections through its fiber optic infrastructure in the Dallas area. In addition to Qwest, the roster of vendors contributing to SCinet reads like a high-tech Who's Who: Nortel Networks, Cisco Systems, Juniper Networks, Marconi Corp., Mitre Corp., Netcom Systems, Foundry Networks, Sun Microsystems, MCI and GST Telecom. SCinet provided connectivity to the nation's leading networks, including the U.S. Department of Energy's Energy Sciences Network, Internet2/Abilene, the National Science Foundation's vBNS, and others.

For the first time in the history of the SC conference, SCinet also provided wireless networking capability throughout the Dallas Convention Center. This allowed conference attendees to connect with SCinet from anywhere on the show floor or in any of the technical program areas. It wasn't unusual to see attendees walking the corridors toward various technical sessions, all the while watching another live session on the laptop computer they were carrying.

The 2000 conference also sponsored a competition for high-bandwidth applications, each of which was expected to completely saturate the network's capability as data from high-performance computing. When the results were announced, NERSC was well represented in the winner's circle. For details, see the following paragraphs.

### **NERSC Team Wins the Battle of the Bandwidth at SC2000 Conference**

A NERSC team of Wes Bethel, Brian Tierney, Jason Lee and Dan Gunter demonstrating Visapult, Wes Bethel's prototype application and framework for remote visualization of terascale datasets, won the "Fastest and Fattest" category for overall best performance in the SC2000 Network Challenge for Bandwidth-Intensive Applications. The team overcame problems with the conference network and equipment in winning the award in the first competition of its kind.

To run their entry, the team used a server at Berkeley Lab running DPSS (the Distributed Parallel Storage System developed by Brian Tierney and others), client machines on the SC2000 show floor running Visapult for remotely visualizing an 80 gigabit data set running on an eight-CPU SGI Origin computer in the ACSI booth, and `dpss_get`, a high-speed parallel file transfer application running on an eight-node Linux cluster in Argonne National Lab's booth at the conference. The Visapult team recorded a peak performance level of 1.48 gigabits per second over a five-second period. Their 60-minute average throughput was 596 megabits per second. Overall, the team transferred more than 2 terabits of data in 60 minutes from LBNL to the SC2000 show floor in Dallas.

Two other NERSC staff members provided key support to another winning team in the Network Challenge. Taking top honors as the "Hottest Infrastructure" application was a team representing the University of Southern California/Information Sciences Institute and Argonne and Lawrence Livermore national laboratories, with help from Berkeley Lab's Arie Shoshani and Alex Sim, running "A Data Management Infrastructure for Climate Modeling Research." The team demonstrated its infrastructure for secure, high-performance data transfer and replication for large-scale climate modeling data sets and achieved a peak performance level of 1.03 gigabits per second.

## TECHNOLOGY TRANSFER

*Goal: Export knowledge, experience, and technology developed at NERSC, particularly to and within NERSC client sites.*

In pursuing our goal of accelerating the rate of scientific discovery, the NERSC staff realizes that the exchange of knowledge and experience is a significant component of our effort. The NERSC staff has compiled an impressive track record in developing new tools, refining existing ones and working actively with specialized groups within the high performance computing community. And just as NERSC staff members share their expertise through technical presentations, papers and articles, we also regularly host visits by representatives of other centers to compare note, build relationships and look for areas of possible collaborations. Here is a partial list of some of our efforts in this field.

### **NERSC-Developed ESP Being Adopted by Other Centers**

Effective System Performance (ESP), a NERSC-developed benchmark that measures in a real-world operational environment, is catching on. After being a featured topic at two consecutive SC conferences (SC99 and SC2000), ESP is now being incorporated into system evaluation tests used at the National Center for Atmospheric Research and Lawrence Livermore National Laboratory. At LLNL, ESP will be a key test for evaluating the new ASCI system. For more information about ESP, see the Wise Technology Integration section of this report.

### **NERSC Provides Tools for High Performance Computing**

In 2000, NERSC was the source of several software packages aimed at improving high-performance computing in both computer science and computational science. In conjunction with the SC2000 Conference in Dallas, NERSC released the “Berkeley Lab AMR” (Adaptive Mesh Refinement) library and the Berkeley Lab VIA Software package of M-VIA and MVICH.

Berkeley Lab AMR is a comprehensive library of adaptive mesh refinement (AMR) software and documentation. AMR serves as a “numerical microscope,” allowing researchers to “zoom in” on the specific regions of a problem that are most important to its solution. Rather than requiring that the whole calculation have the same spatial resolution, AMR allows different resolution in different regions of the problem. Areas of interest are covered with a finer mesh than the surrounding regions; for time-dependent problems, the finer meshes are also advanced with a smaller time step. Not having to perform the entire calculation at the finest resolution allows scientists to make the most of available computer resources, so that they can then solve bigger, harder problems.

The tools are general enough that most computational problems involving partial differential equations can potentially benefit from their use. Now, these tools, including visualization software, and extensive documentation supporting their application, have been produced as a CD for general distribution. Berkeley Lab AMR is unique among many AMR codes because of its adaptability to a wide range of applications. Scalable parallelism, and an object-oriented

approach have been built into the design from the very beginning, and assure the flexibility and high performance across multiple platforms.

Also released on CD was the Berkeley Lab VIA Software package of M-VIA and MVICH, Virtual Interface Architecture software for low-latency, high-bandwidth, inter-process communication. NERSC's Future Technologies Group developed the software to allow very high performance over cluster networks. M-VIA is a modular version of VIA (Virtual Interface Architecture), an industry standard, high performance communication interface for system area networks. First released in late 1998, M-VIA was the first implementation of VIA for Linux. To support the goal of portability to new network interfaces, M-VIA was designed to be highly modular, hence the name M-VIA, for Modular VIA. Using M-VIA and a high performance network, two computers in a cluster can exchange data at a rate about 1,000 times faster than the fastest DSL connection.

The second piece of software on the CD is MVICH, an implementation of MPI (Message Passing Interface) for VIA. It is built on the MPICH implementation of MPI from Argonne National Laboratory. The MPICH code is organized in layers in which the upper levels implement general functionality, independent of the underlying communication mechanism. MPICH abstracts the communication layer through the Abstract Device Interface (ADI), where all point-to-point communication takes place. MVICH implements the ADI layer for VIA.

Andrew Canning, a member of NERSC's Scientific Computing Group, was the lead author of a paper describing a parallel version of one of the most accurate first principles models for determining structural, electronic and magnetic properties of crystals and surfaces. Called FLAPW (for full-potential linearized augmented plane-wave), the method is one of the most accurate and widely used methods for this type of research. The new version allows researchers to take advantage of the world's most powerful supercomputers in modeling the electronic structure of materials. Massively parallel computers, for which the code was developed, allow users to "scale up" their simulations from a small number of atoms to hundreds or thousands of atoms. It was thought in the past that the FLAPW method could not be parallelized efficiently but through a unique parallelization scheme that minimized communications, and the use of efficient parallel eigensolvers from the ACTS toolkit, Andrew Canning was able to develop a parallel version of the code running efficiently on up to 1,000 processors.

This code is now being used by DOE researchers, such as Prof. Art Freeman's group at Northwestern (where the original serial code came from) and Prof. Bruce Harmon's group at DOE's Ames Laboratory in Iowa.

This code allows new physics to be studied that requires the study of larger systems than was not possible with serial codes such as nanostructures, defects, impurities, etc. Understanding how even the smallest impurities can affect the overall performance of a material, such as silicon used in computer chips, is critical to the manufacture and development of technologically important materials. Being able to accurately model such materials at the atomic scale can give researchers insight into how the electronic structures of materials are affected by even tiny defects. The

authors used their version to model a system with 686 silicon atoms and 343 palladium atoms using 512 processors on NERSC's Cray T3E supercomputer. Their version of FLAPW is designed to run on massively parallel supercomputers built by Cray, IBM and other manufacturers.

This work was done in collaboration with Prof. Freeman's group at Northwestern University and Wolfgang Mannstadt at the University of Marburg, Germany.

### **NERSC Introduces HSI — The HPSS Interface Utility**

In 1999, NERSC began working with Mike Gleicher of Gleicher Enterprises LLC to develop HSI, an interface to the High Performance Storage System (HPSS). HSI is intended to be more flexible and user-friendly, and to offer a greater range of commands, than the pftp/ftp interface to HPSS. The NERSC HPSS used HIPPI IPI3 capable disk, and the HIPPI IPI-3 driver was not available on newer Cray computers (J90s and T3E). The purpose of the project was to develop a user-level implementation of the third party HIPPI IPI-3 capability, using the standard HIPPI driver. The Protocol is implemented in a portable library which can be linked with HPSS client applications such as HSI and PFTP.

HSI offers different ways to do any particular file movement. Some of its commands look very similar to ftp file transfer commands; others resemble standard Unix file handling commands; still others resemble those from the CFS utility. NERSC deployed HSI in early 2000 and since then, it has been adopted on two clusters at Berkeley Lab and by other centers. Other sites using the HSI client include Lawrence Livermore National Laboratory, Oak Ridge National Laboratory, the San Diego Supercomputer Center, Caltech, Maui High Performance Computing Center, the University of Indiana, the National Climatic Data Center, the University of Maryland and BAE Systems.

### **ORNL, NERSC Representatives Meet to Share Ideas, Opportunities to Collaborate**

Four staff members from Oak Ridge National Laboratory's Computer Science and Mathematics (CSM) Division met with their counterparts in NERSC in July, comparing notes and discussing potential areas for future collaboration. Attending from ORNL were Buddy Bland (director of operations), Randy Buries (storage systems lead), Richard Alexander (operations systems lead) and Trey White (user services). After hearing an overview of the computing centers at ORNL and Berkeley Lab, the visitors toured the Lab's new facility under construction in downtown Oakland. In the afternoon, the visitors met in small groups with NERSC staff members to talk about specific practices, problems and potential areas for further collaboration. The group discussing user services covered such areas as how security information is collected from users, what software is needed by users and whether the same software could be installed at all of DOE's computing sites, making it easier for users. The scientific computing group discussed the various libraries used at the centers and the ACTS Toolkit, a collection of DOE-developed software supported by NERSC.

The storage systems group focused on HPSS (High Performance Storage System), which is the focus of a joint NERSC-ORNL testbed called PROBE. One issue in particular was how to increase the bandwidth between ORNL and LBNL storage systems, each of which have high bandwidth internally but much slower external connections. At a strategic level, there is common interest between the two centers in improving the system management that exists on HPC systems and in the long-term future for HPSS/storage. In system management, both ORNL and NERSC staff have had success in doing focused, value-added research and development for system functionality in limited but key areas. For HPSS, both sites are invested in this system, there is nothing that is now clearly superior, and both are interested in how to make it widely viable in the future.

### **NASA Team Visits NERSC as Part of International 'Best Practices' Survey**

Five NASA researchers from around the country met with NERSC representatives in August as part of a two-month survey of 24 computing organizations in the U.S. and Europe. The visit was part of NASA's "Best Practices Data Center Study," which is designed to gather information to help manage the data and information management system for NASA's Earth Observing System (EOS). EOS is a series of space-based missions designed to measure key parameters needed to understand global climate change. The project is expected to generate terabytes of data by 2004.

The visit also included a tour of NERSC's computing facilities. Among the topics NASA was interested in were HPSS mass storage, distributed computing environments, computing resources, archival storage and data migration. The visitors were also interested in the related issues of staffing levels, staffing qualifications, and how staffing levels for various NERSC "data center" functions are established. "Staffing is important to NASA in this Benchmark Study -- we are trying to get a better handle on it by visiting various data centers and hopefully getting some good ideas," the team wrote in requesting the meeting.

The NASA representatives were Ron Holland of the NASA Earth Science Data and Information Systems Project in Maryland, Ron Weaver from the National Snow Ice Data Center in Colorado, Tom Kalvelage of the EROS Data Center in South Dakota, Arthur "Bud" Booth from SGT Inc. in Maryland and John Moses from the Computer Sciences Corporation in Maryland.

### **NERSC, SDSC Representatives Meet, Discuss Areas of Common Interest**

NERSC hosted a day-long meeting in July with representatives of the San Diego Supercomputer Center to discuss technical areas of common interest, as well as help staff at the two centers get to know each other as a means to improving communications in the future. The July 6 meeting grew out of conversations held at the NPACI All-Hands Meeting in January. Both centers are currently home to large IBM SP supercomputers and High Performance Storage Systems (HPSS).

Areas of potential collaboration include SP administration and management, Grid tools deployment, clusters, performance characterization and modeling, training and DOE's ACTS

Toolkit. NERSC also agreed to provide SDSC information on the Best Value Source Selection criteria used by NERSC in procuring its IBM system. The process looks not only at the price of a system, but also its performance level to determine the best overall value of various competing systems.

### **NERSC Makes Real-time Network Analysis Tool “pipechar” Available**

A new and easy-to-use tool for analyzing and monitoring the network was developed by Jin Guojun of the Data Intensive Distributed Computing Group in NERSC’s Distributed Systems Department. The tool, called pipechar, is a sub-service of the Network Character Service daemon (NCSD) and has been extracted as an individual tool paired with netest for network analysis and monitoring. According to Jin, pipechar is a simple tool which users can run themselves from their desktop computers. This capability of analyzing network traffic is important in order for distributed applications to fully utilize high-speed WANs in the near future, and for doing network layer congestion manipulation and Quality of Service (QoS).

### **NERSC Hosts Workshop on ACTS Toolkit for Grad Students, Post-Docs**

In September, NERSC hosted a three-day workshop to familiarize graduate students and post-docs with the ACTS, or Advanced Computational Testing and Simulation, Toolkit. The ACTS Toolkit is a set of mostly DOE-developed software tools that make it easier for programmers to write high-performance scientific applications for parallel computers. The workshop drew about 70 applications for the 30 available slots.

The aim of the workshop, entitled “How can ACTS work for you?” is to familiarize students in various scientific disciplines with the computational science tools available to help them in their research. The tools in the ACTS Toolkit provide solutions for numerical problems, scientific data representations, data manipulation, visualization, program execution and distributed computing.

Over the past year, NERSC employees Osni Marques and Tony Drummond have organized several events and made presentations around the country to raise HPC community awareness of the ACTS Toolkit. For more information visit the ACTS website at <<http://acts.nersc.gov>>.

### **Active Members of the HPC Community**

NERSC staff also share their expertise through presentations, publications and by giving workshops and tutorials, reaching both specialized and general audiences. For example, in February 2000, NERSC hosted the Bay Area Scientific Computing Workshop. The workshop, which drew more than 100 scientists from around Northern California, was an informal gathering to encourage the interaction and collaboration of Bay Area researchers in the field of scientific computing. The organizers, members of NERSC’s Scientific Computing Group, were particularly interested in providing a forum for young or beginning researchers in the field to talk about their work.

Here are some of the other ways in which NERSC staff contributed to the exchange of ideas in the HPC community during the year 2000. The following lists are representative, but not necessarily exhaustive:

### **Invited Talks**

- “Fast million atom pseudopotential calculations using SLCBB method,” Lin-Wang Wang, American Physical Society annual meeting, Minneapolis, Minn., March 2000.
- “Modeling Electrostatic Effects in Soft Matter,” Niels Gronbech-Jensen, Invited Visiting Professor, Centro Atomico Bariloche, Bariloche, Argentina, March 2000.
- “Large Scale Materials Science Simulations,” Andrew Canning, DOE Computational Science Graduate Fellowship Conference, Pleasanton, CA, March 2000.
- “Using Accurate Arithmetics to Improve Numerical Reproducibility and Stability in Parallel Applications,” Chris Ding, Data Assimilation Office, NASA Goddard Space Flight Center, Greenbelt, Md., March 2000
- “Data Organization and I/O in Modular Ocean Model,” Chris Ding, DOE BER/ESD CCPP PI Science Meeting, Bethesda, Md., March 2000.
- “Are the Digits of Pi Random?” David Bailey, Center for Computing Sciences, April 2000.
- “Large Scale parallel materials science simulations,” Andrew Canning, University of California at Davis, CA, April 2000.
- “The ordering problem for sparse matrix factorizations,” Esmond G. Ng, UC Davis, Davis, Calif., April 2000.
- “Single-particle Reconstruction Using Supercomputing,” Esmond G. Ng. Single-particle Reconstruction Workshop at the National Center for Macromolecular Imaging, Houston, Texas, April 2000.
- “Support Vector Machines and its Application in Protein Fold Predictions,” Chris Ding, Computational/Structural Biology Seminar, Lawrence Berkeley National Laboratory, April 2000.
- “A Probabilistic Model for LSI/SVD in Information Retrieval,” Chris Ding, Scientific Computing & Computational Mathematics Seminar, Stanford University, Palo Alto, May 2000.
- “What’s New in Linear Equations Solvers,” Esmond G. Ng, Workshop on Computational Challenges in Atomic and Molecular Physics, Institute for Theoretical Atomic and Molecular Physics, Harvard University, Mass., May 2000.
- “SNAP: Supernova / Acceleration Probe, An Experiment to Measure the Properties of the Accelerating Universe,” Peter Nugent, Arecibo Observatory, Puerto Rico, May 2000.

“Challenges of Future High-End Computing,” (invited plenary talk), David Bailey, SP World Conference, San Diego, June 2000.

“Mid Range Computing at LBNL,” Bill Kramer, LBNL Computing and Communications Services Advisory Council Meeting, July 2000.

“Phase-locking of Vortex Lattice Motion,” Niels Gronbech-Jensen, Future Perspectives of Superconducting Josephson Devices; Physics and Applications of Multi-Junctions Superconducting Josephson devices, Acquafredda di Maratea, Italy, July 2000.

“New Ideas For Solving Sparse Nonsymmetric Linear Systems,” Esmond G. Ng, Weierstrass Institute for Applied Analysis and Stochastics, Berlin, Germany, August 2000.

“The Ordering Problem for Sparse Matrix Factorization,” Esmond G. Ng, the 16th IMACS World Congress on Scientific Computation, Modeling and Applied Mathematics, Lausanne, Switzerland, August 2000.

“Visapult - A Prototype Remote and Distributed Visualization Application and Framework,” Wes Bethel, Gigabit Networking Workshop sponsored by NASA Ames Lab and the NASA Research and Education Network, Moffet Field, Calif., August 2000.

“New Ideas For Solving Sparse Nonsymmetric Linear Systems,” Esmond G. Ng, ETH (Swiss Federal Institute of Technology), Zurich, Switzerland, August 2000.

“Emerging Technologies for Remote Parallel Visualization,” Wes Bethel, First SIAM Conference on Computational Science and Engineering, Washington, D.C., September 2000.

“The NERSC Effective System Performance Test,” Bill Kramer, IDC Workshop on High Performance Computing, September 2000.

“Predicting dopant implant profiles in semiconductors,” Niels Gronbech-Jensen, Department of Physics, San Diego State University, San Diego, Calif., September 2000.

“Opinions on Future HPC Systems,” Bill Kramer, Cray Inc. Mission Critical Customer Advisory Board Meeting, Richmond, VA, September 2000.

“Disordered Particle Systems with Long-Range Interactions,” Niels Gronbech-Jensen, Department of Mathematics, University of North Carolina, Chapel Hill, N.C., October 2000.

“Scientific Computing at NERSC”, Esmond G. Ng, Colloquium, Department of Computer Science, University of Kentucky, Lexington, Ky. October 2000.

“Linear Algebra and Scientific Discoveries,” Esmond G. Ng, SIAM Applied Linear Algebra, Raleigh, N.C., October 2000.

“A Probabilistic Model for LSI/SVD in Information Retrieval,” Chris Ding, SIAM Computational IR Workshop, October 2000.

“Million atom electronic structure calculations of nanostructures,” Lin-Wang Wang, Physics Department seminar, UC Davis, Davis, Calif. October 2000.

“Human Genome Project Overview,” Manfred Zorn, Second State and Federal Justices Relationship Conference, Tucson, Ariz, October 2000.

“Linear Algebra, What’s it Worth?” Horst Simon (presented by Esmond Ng), SIAM (Society for Industrial and Applied Mathematics) Conference on Applied Linear Algebra, Raleigh, N.C. October 2000.

“Large Scale Parallel Excited State Calculations at NERSC,” Andrew Canning, Workshop on excited state properties and response functions for materials, Minneapolis, November 2000

“Dimensionality Reduction in Information Retrieval and Filtering,” Chris Ding, Colloquia, Dept of Computer Science, UC Davis, November 2000.

“Recent Developments in Sparse Unsymmetric Factorization: Parallelism and Fill-in Reducing Ordering,” Xiaoye Sherry Li, Computational Sciences and Mathematics Research Department, Sandia National Laboratories’ Livermore site, November 2000.

“Cosmic Microwave Background Radiation (CMB) data analysis,” Julian Borrill, Computer Science Department, UC Davis, Davis, Calif., November 2000.

### **Conference Presentations and Proceedings**

“XXth Moriond Workshop: Energy Densities in the Universe,” Julian Borrill, Moriond, January 2000.

“Interactive Supercomputing,” Parry Husbands, NAM Granville-Browne Presentation at the AMS Joint Meetings, Washington D.C. January 2000.

“Using NetLogger for Distributed Systems Performance Analysis of the BaBar Data Analysis System,” Brian Tierney, Dan Gunter, J. Becla, B. Jacobsen, David Quarrie, Proceedings of Computers in High Energy Physics 2000 (CHEP 2000), February 2000.

“The CDF Run II Data Catalog and Data Access Modules,” Paolo Calafiura, CHEP 2000, International Conference on Computing in High Energy and Nuclear Physics, Padova, Italy, February 2000.

“New and current algorithms for solving the Kohn-Sham equations of density functional theory: SciDac Thomas-Fermi charge mixing for obtaining self-consistency in density functional calculation,” David Raczowski, American Physical Society annual meeting, Minneapolis, Minn., March 2000.

“Parallelization of the FLAPW method and comparison to the PPW method,” Andrew Canning, American Physical Society March Meeting, Minneapolis, Minn., March 2000.

- “Electronic structure calculations using large scale parallel computers,” Lin-Wang Wang, Material Research Society (MRS) spring meeting, San Francisco, Calif., April 2000.
- “Implementing a Visualization Tool for Adaptive Mesh Refinement Data using VTK,” Terry J. Ligocki, Visualization Development Environments 2000, Princeton Plasma Physics Lab., Princeton, NJ, April 2000.
- “Using Accurate Arithmetics to Improve Numerical Reproducibility and Stability in Parallel Applications,” Yun Helen He and Chris. H.Q. Ding, presented at International Conference on Supercomputing (ICS’00) in Santa Fe, New Mexico, May 2000; and at the Ninth Workshop on the Use of High Performance Computing in Meteorology: Developments in Teracomputing, European Centre for Medium-Range Weather Forecasts, Reading, U.K., November 2000.
- “System Utilization Benchmark on the Cray T3E and IBM SP,” Adrian Wong, Leonid Oliker, William Kramer, Teresa Kaltz, David Bailey, 14th International Parallel and Distributed Processing Symposium (IPDPS), Cancun, Mexico, May 2000.
- “System Utilization Benchmark on the Cray T3E and IBM SP,” Adrian T. Wong, Leonid Oliker, William T. C. Kramer, Teresa L. Kaltz and David H. Bailey, Fifth Workshop on Job Scheduling, May 2000.
- “Remote Control for Videoconferencing,” Deb Agarwal and Marcia Perry, Proceedings of the 11th International Conference of the Information Resources Management Association, Anchorage, Alaska, May 2000.
- “Remote Control for Videoconferencing,” Marcia Perry, DOE2000 Review, Argonne National Laboratory, Ill., May 2000.
- “Using Accurate Arithmetics to Improve Numerical Reproducibility and Stability in Parallel Applications,” Helen. He and Chris Ding, Proceedings of International Conference on Supercomputing (ICS’00), May 2000.
- “WE1997/WE2000,” D. W. Vasco, L. Johnson, R. J. Pulliam, P. Earle and O. A. Marques (presented by D. W. Vasco), Reference Earth Model Workshop, Harvard, Cambridge, Mass., May 2000.
- “A higher-order embedded boundary method for time-dependent simulation of hyperbolic conservation laws,” David Modiano and Phil Colella Proceedings of the ASME 2000 Fluids Engineering Division Summer Meeting, Boston, Mass., June 2000.
- “Feature Based Visualization of Large Scale Geophysical Data,” Qing Yang and Bahram Parvin, IEEE Conference on Computer Vision and Pattern Recognition, June 2000.
- “A New Regularized Approach for Contour Morphing,” Ge Cong and Bahram Parvin, IEEE Conference on Computer Vision and Pattern Recognition, June 2000.
- “Interactive Supercomputing,” Parry Husbands, CAARMS6, Baltimore, Md., June 2000.

- “Accurate Singular Values with dqds,” Osni A. Marques, International Workshop on the Accurate Solution of Eigenvalue Problems, Hagen, Germany, July 2000.
- “9th Marcel Grossman Meeting,” Julian Borrill, Universita de Roma La Sapienza, Rome, July 2000.
- “ESO Astrophysics Symposium: Mining The Sky,” Julian Borrill, Max Planck Institute, Garching, Germany, August 2000.
- “Computacao de Alto Desempenho: Calculo de Autovalores e Ferramentas de Programacao,” Osni A. Marques, minicourse taught in Portuguese, PUC/RJ, Rio de Janeiro, Brazil, August 2000.
- “Effective System Performance Benchmark,” Adrian Wong, Leonid Oliker, William Kramer, Teresa Kaltz, Therese Enright, David Bailey, presented at SP World, June 2000, San Diego, Calif., June 2000, and at SP Scicom 2, San Diego, Calif., August 2000.
- “NetLogger: A Toolkit for Distributed System Performance Analysis,” Dan Gunter, Brian Tierney, Brian Crowley, M. Holding, Jason Lee, Proceedings of the IEEE Mascots 2000 Conference (Mascots 2000), August 2000.
- “A Monitoring Sensor Management System for Grid Environments,” Brian Tierney, Brian Crowley, Dan Gunter, M. Holding, Jason Lee, Mary Thompson, Proceedings of the IEEE High Performance Distributed Computing conference ( HPDC-9 ), Pittsburgh, Pa., August 2000.
- “Parallel Conjugate Gradient: Effects of Ordering Strategies, Programming Paradigms, and Architectural Platforms,” Sherry Li, Leonid Oliker, G. Heber and R. Biswas, Proceedings of the 13th International Conference on Parallel and Distributed Computing Systems, August 2000.
- “Parallelization of the FLAPW method and applications to large systems,” Andrew Canning, psik2000 conference, Schwaebisch Gmuend, Germany, August 2000.
- “The Cactus Code: A Problem Solving Environment for the Grid,” Gabrielle Allen, Werner Benger, Tom Goodale, Hans-Christian Hege, Gerd Lanfermann, André Merzky, Thomas Radke, Edward Seidel and John Shalf, Proceedings of the Ninth IEEE International Symposium on High Performance Distributed Computing (HPDC9), Pittsburgh, August 2000.
- “Visapult - A Prototype Remote and Distributed Visualization Application and Framework,” Wes Bethel, Proceedings of Siggraph 2000, New Orleans La., August 2000.
- “NASA Advanced Information Systems Research Projects,” Julian Borrill, South West Research Institute, Boulder, September 2000.
- “Singular Features from Sea Surface Temperature Data,” Qing Yang and Bahram Parvin, International Conference On Pattern Recognition, September 2000.
- “Shape from Cross Sectional Contours,” Ge Cong and Bahram Parvin, International Conference On Pattern Recognition, September 2000.

“PDSF Site Report,” Iwona Sakrejda, HEPIX-HEPNET 2000 Conference at Jefferson Lab, Newport News, Va., October 2000.

“Initial Results of the CD-1 Reliable Multicast Experiment,” Deb Agarwal with R. Stead, B. Coan, J. E. Burns, N. Shah, N. Kyriakopoulos, Proceedings of the GCI Workshop 2000, Vienna, Austria, October 2000.

“Symmetric Minimum Priority Orderings for Sparse Unsymmetric Factorization,” Xiayoe Sherry Li, Patrick Amestoy and Esmond Ng, SIAM (Society for Industrial and Applied Mathematics) Conference on Applied Linear Algebra, Raleigh, N.C. October 2000.

“A Latency Tolerant Scheme for Parallel Preconditioning with Incomplete Factors,” Esmond Ng, Padma Raghavan and Keita Teranishi (presenter), SIAM Conference on Applied Linear Algebra, Raleigh, N.C. October 2000.

“A Probabilistic Model for LSI/SVD in Information Retrieval,” Chris Ding, Workshop on Computational Informational Retrieval in conjunction with the SIAM Conference on Applied Linear Algebra, Raleigh, N.C. October 2000.

“LAPACK’s Singular Value Program,” Beresford N. Parlett and Osni A. Marques, Seventh SIAM Conference on Applied Linear Algebra, Raleigh, N.C. October 2000.

“On the Use of the Singular Value Decomposition for Text Retrieval,” Parry Husbands, Workshop on Computational Informational Retrieval in conjunction with the SIAM Conference on Applied Linear Algebra, Raleigh, N.C. October 2000.

“A Framework for Visualizing Hierarchical Computations,” Terry J. Ligocki, Brian Van Straalen, John M. Shalf, Gunther H. Weber, Bernd Hamann, NSF/DOE Lake Tahoe Workshop on Hierarchical Approximation and Geometrical Methods for Scientific Visualization 2000, Granlibakken Conference Center, Tahoe City, Calif., October 2000.

“Extraction of Crack-free Isosurfaces from Adaptive Mesh Refinement Data,” Gunther H. Weber, Oliver Kreylos, Terry J. Ligocki, John M. Shalf, Hans Hagen, Bernd Hamann, NSF/DOE Lake Tahoe Workshop on Hierarchical Approximation and Geometrical Methods for Scientific Visualization 2000, Granlibakken Conference Center, Tahoe City, Calif., October 2000.

“On the Use of Singular Value Decomposition for Text Retrieval,” Parry Husbands, Horst Simon and Chris Ding, Proceedings of 1st SIAM Computational Information Retrieval Workshop, Raleigh, N.C., October 2000.

“Numerical Reproducibility on Distributed Platforms: An Accurate Arithmetics Approach,” Yun Helen He and Chris H.Q. Ding, Research Gem poster at SC2000, Dallas, TX, November 2000.

“BioSig: A Bioinformatic System for Studying the Mechanism of Inter-cell Signaling,” Bahram Parvin, Ge Cong, Gerald Fontenay, R. Henshall and Mary Helen Barcellos-Hoff, IEEE International Symposium on Bioinformatics and Biomedical Engineering, November 2000.

“ESP: A System Utilization Benchmark,” Adrian T. Wong, Leonid Oliker, William T. C. Kramer, Teresa L. Kaltz and David H. Bailey, Proceedings of the IEEE SC2000 Conference, Dallas, November 2000.

“Using High-Speed WANs and Network Data Caches to Enable Remote and Distributed Visualization,” Wes Bethel, Brian Tierney, Jason Lee, Dan Gunter and Stephen Lau, IEEE SC2000 Conference, Dallas, November 2000.

“Stochastic optimal prediction with application to Averaged Euler equations,” John Bell, Alexandre Chorin and William Crutchfield, Proceedings of the 7<sup>th</sup> National Conference on Computational Fluid Mechanics, Pingtung, Taiwan, 2000.

“Extending OLAP Querying to External Object Databases,” T. Pedersen, Arie Shoshani, Jumin Gu, C.S. Jensen, 9th International Conference on Information and Knowledge Management (CKIM’00), 2000.

“Coordinating Simultaneous Caching of File Bundles from Tertiary Storage,” Arie Shoshani, Alex Sim, Luis M. Bernardo, Henrik Nordberg, 12th International Conference on Scientific and Statistical Database Management (SSDBM’00).

Poster on NADIMS, a complex metadatabase application using OPM, to Carl Anderson, Mary Anderson and Arie Shoshani, Scientific Data Management Group, American Geophysical Union’s Fall Meeting, San Francisco, December 2000.

“Supernova Evolution and its Impact on Cosmological Measurements,” Peter Nugent, D. Kasen, E. Baron, D. Branch, American Astronomical Society Meeting proceedings, December 2000.

### **Workshops and Tutorials**

“The ACTS Toolkit: An Overview,” Osni A. Marques, NPACI AHM2000, San Diego, Calif., February 2000.

Department of Energy Computer Graphics Forum, Wes Bethel, Technical Chair, Santa Fe, N.M., May 2000.

“The ACTS Toolkit: What can it do for you?” Osni A. Marques, NERSC Users Group Meeting, Oak Ridge, Tenn., June 2000.

WOMPAT2000: Workshop on OpenMP Applications and Tools, Helen He: San Diego Supercomputer Center, San Diego, Calif., July 2000.

“Tutorial on the ACTS Toolkit,” Tony Drummond, San Diego Supercomputing Center, Parallel Programming Summer Institute, La Jolla, Calif., July 2000.

“Computational Ocean Modeling,” Yun Helen He and Chris H.Q. Ding, DOE Computational Science Graduate Fellow (CSGF) Workshop, Lawrence Berkeley National Laboratory, Berkeley, Calif., July 2000.

“SNAP Scientific Working Group on Core Collapse Supernovae,” Peter Nugent, organizer, Berkeley, Calif., August 2000.

“A Cache-based Data Intensive Distributed Computing Architecture for Grid Applications,” Brian Tierney, William Johnston, Jason Lee, CERN School of Computing, September 2000.

“The ACTS Toolkit Workshop: How can ACTS work for you?” Osni Marques and Tony Drummond, Berkeley, Calif., September 2000.

“A Dynamic State Machine for the Atlas Software Framework,” Paolo Calafiura, ACAT 2000, VII International Workshop on Advanced Computing and Analysis Techniques in Physics Research, Fermilab, Illinois, October 2000.

“Introduction to PDSF,” Iwona Sakrejda, STAR Workshop, Lawrence Berkeley National Laboratory, November 2000.

“Computational Biology and High Performance Computing,” Manfred Zorn, Sylvia Spengler, Inna Dubchak and Craig Stewart, SC2000 Full-Day Tutorial, Dallas, Texas, November 2000.

“Performance Tuning and Analysis for Grid Applications,” Brian Tierney, Dan Gunter, Rich Wolski and Martin Swamy, SC2000 Half-Day Tutorial, Dallas, Texas, November 2000.

“The ACTS Toolkit Development,” Tony Drummond and Osni Marques, SC2000, Dallas, Texas, November 2000.

“Computational Biology and High Performance Computing,” Manfred Zorn, Technical University, Dresden, Germany, November 2000.

“Code Development for Materials Science,” Andrew Canning, Workshop on a topical center for Materials Science, Oak Ridge National Lab, Oak Ridge, November 2000.

### Published papers and articles

“Integer Relation Detection,” David H. Bailey, *Computing in Science and Engineering*, January/February 2000.

“Small Scale Processes and Entrainment in a Stratocumulus Marine Boundary Layer,” David E. Stevens, John B. Bell, Ann S. Almgren, Vincent E. Beckner and Charles A. Rendleman, *Journal of Atmospheric Sciences*, February 2000.

“A Flat Universe from High-Resolution Maps of the Cosmic Microwave Background Radiation”, Julian Borrill, P. de Bernardis et al., *Nature*, April 2000.

“Using Accurate Arithmetics to Improve Numerical Reproducibility and Stability in Parallel Applications,” Yun Helen He and Chris H.Q. Ding, *Proceedings of International Conference on Supercomputing (ICS'00)*, May 2000.

“Visualization Dot Com,” Wes Bethel, *IEEE Computer Graphics and Applications*, May/June 2000.

“Shape Recovery from Equal Thickness Contours,” Ge Cong and Bahram Parvin, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, August 2000.

“Model Based Segmentation of Nuclei,” Ge Cong and Bahram Parvin, *Journal of Pattern Recognition*, August 2000.

“Modelling surface growth in Ion Beam Assisted Deposition with rate equations,” Jussi K. Sillanpaa, Ismo T. Koponen and Niels Gronbech-Jensen, *Proceedings from Materials Research Society Meeting*, Boston, November 2000.

“Experimental Mathematics: Recent Developments and Future Outlook”, David H. Bailey and Jonathan M. Borwein, *Mathematics Unlimited - 2001 and Beyond*, Springer-Verlag, November 2000.

“The Effect of Stoichiometry on Vortex Flame Interactions,” John Bell, Nancy Brown, Marcus Day, Michael Frenklach, Joseph. Grcar and Shaheen Tonse, *Proceedings of the Combustion Institute*, 2000.

“Scaling and Efficiency of PRISM in Adaptive Simulations of Turbulent Premixed Flames,” John Bell, Nancy Brown, Marcus Day, Michael Frenklach, Joseph. Grcar, Rick Propp and Shaheen Tonse, *Proceedings of the Combustion Institute*, 2000.

“Numerical Simulation of Laminar Reacting Flows with Complex Chemistry,” Marcus S. Day and John B. Bell, *Combustion Theory Modeling*, 2000.

“A numerical model for trickle-bed reactors,” Rick Propp, Phil Colella, William Crutchfield and Marcus Day, *Journal of Computational Physics*, 2000.

“Parallelization of Structured, Hierarchical Adaptive Mesh Refinement Algorithms,” Charles A. Rendleman, Vince E. Beckner, Michael Lijewski, William Y. Crutchfield, John B. Bell, *Computing and Visualization in Science*, Volume 3, 2000.

“Parallelization of Structured, Hierarchical Adaptive Mesh Refinement Algorithms,” Charles A. Rendleman, Vince E. Beckner, Michael Lijewski, William Y. Crutchfield, John B. Bell, *Computing and Visualization in Science*, Volume 3, 2000.

“Access Coordination of Tertiary Storage for High Energy Physics Application,” L. M. Bernardo, A. Shoshani, A. Sim, H. Nordberg, MSS 2000.

“A Building Block Approach for Determining Low-frequency Normal Modes of Macromolecules,” Osni Marques (with F. Tama, F. X. Gadea and Y.-H. Sanejouand), *Proteins: Structure, Function and Genetics*, 2000.

“An Implementation of the dqds Algorithm (positive case),” Osni Marques with Beresford Parlett, *Linear Algebra and its Applications*, 2000.

“First Estimations of Cosmological Parameters From BOOMERANG”, Julian Borrill, A.E. Lange et al., *Physical Review D*, 2000.

“MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on angular scales of 10 arcminutes to 5 degrees”, Julian Borrill, S. Hanany et al., *Astrophysical Journal Letters*, 2000.

“Constraints on Cosmological Parameters from MAXIMA-1”, Julian Borrill, A. Balbi et al., *Astrophysical Journal Letters*, 2000.

“Asymmetric Beams in Cosmic Microwave Background Anisotropy Experiments,” Julian Borrill, J. H. P. Wu et al., *Astrophysical Journal Supplement*.

“Cosmology from MAXIMA-1, BOOMERANG and COBE/DMR CMB Observations”, Julian Borrill, A. H. Jaffe et al., *Physical Review Letters*, 2000.

“Parallel Empirical Pseudopotential Electronic Structure Calculations for Million Atom Systems,” Andrew Canning, Lin-Wang Wang, A Williamson and A. Zunger, *Journal of Computational Physics*, 2000.

“Structure and energy of the 90 degree partial dislocation in diamond: a combined ab initio and elastic theory analysis,” X. Blase, K. Lin, Andrew Canning, S.G. Louie, D.C. Chrzan, *Physical Review Letters*, 84, 2000.

“Parallelization of the FLAPW Method,” Andrew Canning, W. Mannstadt and A.J. Freeman, *Computer Physics Communications* 130, 2000.

“Multisensor Studies on El Nino-Southern Oscillations and Variabilities in Equatorial Pacific,” X.-H. Yan, Yun Helen He, R. D. Susanto, and W. T. Liu, *Journal of Advanced Marine Sciences and Tech. Society*, 4, 2000.

“An Implementation of the dqds Algorithm (Positive Case),” B. N. Parlett and Osni A. Marques, *Linear Algebra and its Applications*, 2000.

“A Building Block Approach for Determining Low-frequency Normal Modes of Macromolecules,” F. Tama, F. X. Gadea, Y.-H. Sanejouand and Osni A. Marques, *Proteins: Structure, Function and Genetics*, 2000.

“Linear Systems,” Esmond G. Ng and Barry Peyton, *Handbook of Discrete and Combinatorial Mathematics*, ed. Kenneth H. Rosen, CRC Press, 2000.

“Towards a scalable hybrid sparse solver,” Esmond G. Ng and Padma Raghavan, *Concurrency: Pract. Exper.* 2000.

“Statistical mechanics of a discrete nonlinear system,” K. Rasmussen, T. Cretegny, P.G. Kevrekidis, and Niels Gronbech-Jensen, *Physical Review Letters*, 2000.

“Phase-locking of vortex lattices interacting with periodic pinning,” Charles Reichhardt, Richard T. Scalettar, Gergely T. Zimanyi, and Niels Gronbech-Jensen, *Physical Review B*, 2000.

“Interactions between charged rods near salty surfaces,” Rebecca Menes, Niels Gronbech-Jensen, and Philip A. Pincus, *European Physical Journal E*, 2000.

“The effect of splayed pins on vortex creep and critical currents,” C.J. Olson, R.T. Scalettar, G.T. Zimanyi, and Niels Gronbech-Jensen, *Physical Review B*, 2000.

“Charge ordering and long-range interactions in layered transition metal oxides: a quasiclassical continuum study,” Branko Stojkovic, Z.G. Yu, A.L. Chernyshev, A.R. Bishop, A.H. Castro Neto, and Niels Gronbech-Jensen, *Physical Review B*, 2000.

“Collective multi-vortex states in periodic arrays of traps,” Charles Reichhardt and Niels Gronbech-Jensen, *Physical Review Letters*, 2000.

“Shapiro steps in driven vortex lattices interacting with periodic pinning arrays,” C. Reichhardt, R.T. Scalettar, G.T. Zimanyi and Niels Gronbech-Jensen, *Physica C*, 2000.

“Dynamical ordering in the c-axis in 3D driven vortex lattices,” Alejandro B. Kolton, Daniel Dominguez and Niels Gronbech-Jensen, *Physica C*, 2000.

“Vortex pinning and dynamics in layered superconductors with periodic pinning arrays,” Charles Reichhardt, Cynthia J. Olson and Niels Gronbech-Jensen, *Physica C*, 2000.

“Depinning and dynamic phases in driven three-dimensional vortex lattices in anisotropic superconductors,” Cynthia J. Olson and Niels Gronbech-Jensen, *Physica C*, 2000.

“Ordenamiento dinamico en el eje c en superconductores laminares altamente anisotropicos,” Alejandro B. Kolton, Daniel Dominguez, Cynthia J. Olson, and Niels Gronbech-Jensen, *Revista de la Asociacion Fisica Argentina*, 2000.

“Sincronizacion temporal en redes de vortices fuera del equilibrio en medios desordenados,” Alejandro B. Kolton, Daniel Dominguez and Niels Gronbech-Jensen, *Revista de la Asociacion Fisica Argentina*, 2000.

“Cosmological-Model-Parameter Determination from Satellite-Acquired Type Ia and IIP Supernova Data,” S. Podariu, Peter Nugent and B. Ratra, *Astrophysical Journal*, 2000.

“A Strategy for Finding Gravitationally-Lensed Distant Supernovae,” M. Sullivan, R. Ellis, Peter Nugent, I. Smail and P. Madau, *Monthly Notices of the Royal Astronomical Society*, 2000.

“The Rise-Times of High and Low Redshift Type Ia Supernovae Are Consistent,” G. Aldering, R. A. Knop and P. Nugent, *Astronomical Journal*, 2000.

“Metallicity Effects in NLTE Model Atmospheres of Type IA Supernovae,” E. Lentz, E. Baron, D. Branch, P. Hauschildt and Peter Nugent, *Astrophysical Journal*, 2000.

“The acceleration of the Universe: measurements of cosmological parameters from type Ia supernovae,” A. Goobar, S. Perlmutter, G. Aldering, G. Goldhaber, R.A. Knop, Peter Nugent, et al., *Physica Scripta*, volume T, 2000.

“Parallel Empirical pseudopotential electronic structure calculations for million atom systems,” Andrew Canning, Lin-Wang Wang, A. Williamson, and Alex Zunger, *Journal of Computational Physics*, 2000.

“Predicting the size and temperature dependent shapes of precipitates in Al-Zn alloys,” S. Muller, C. Wolverton, L.W. Wang, and A. Zunger, *Acta Mater*, 2000.

“Theoretical interpretation of the experimental electronic structure of lens-shaped self-assembled InAs/GaAs quantum dots,” A.J. Williamson, L.W. Wang, and A. Zunger, *Phys. Rev. B*, 2000.

Book review of “Quantum Wells, Wires and Dots,” by Paul Harrison, Lin-Wang Wang, *Computational Physics Communications*, 2000.

### **Visitors from Across the Country, Around the World**

NERSC regularly hosts visitors from other computing centers and research institutions. These visits range from informal meetings and small group interactions to formal presentations and detailed discussions. Here is a partial list of visitors during 2000.

#### **January**

Edward F. Sowell, Dept. of Computer Science, California State University at Fullerton

Jonathan Emery, LLNL, Livermore, Calif.

Oliver Kreylos, UCD, Davis, Calif.

Iain S. Duff, CERFACS and RAL, Toulouse, France

**February**

Fred G. Gustavson, Mathematical Sciences Department, IBM T.J. Watson Research Center, Yorktown Heights, N.Y.

George Lake, U. of Washington, Astronomy Dept., Seattle, Wash.

Nelson Max, LLNL/ UC Davis, Livermore, Calif.

Carla Maria Dal Sasso Freitas, Federal University of Rio Grande do Sul, Porto Alegre, Rio Grande do Sul, Brazil

Tony Drummond, UCLA, Los Angeles, Calif.

Matthias Bollhoefer, Chemnitz University, Germany, and University of Minnesota

Yann-Herve De Roeck, IFREMER/IRISA, Rennes, France

**March**

Daniel Boley, University of Minnesota and Stanford University

Mike Gleicher, Gleicher Enterprises, San Diego, Calif.

**April**

Adam Arkin, UC Berkeley, Department of Bioengineering and Chemistry

Cleve Ashcraft, Boeing Phantom Works, Seattle, Wash.

**May**

David Raczkowski, UC Davis, Davis, Calif.

Kwan-Liu Ma, UC Davis, Davis, Calif.

Greg Wilson, Nevex Software Technologies, San Mateo, Calif.

David Blackston, UC Berkeley.

Ed Seidel and Dr. Gabrielle Allen, Max-Planck Institute for Gravitational Physics, Potsdam, Germany

**June**

Woo Sun Yang, LANL; UCLA-IGPP Ching-Yao Fong, UC Davis, Davis, Calif.

Justin Wan, Stanford University, Palo Alto, Calif.

Achilleas S. Frangakis, Max-Planck Institute for Biochemistry, Martinsried, Germany

Rasmus Munk Larsen, Stanford University, Palo Alto, Calif.

Joe Grcar, Sandia National Laboratories, Livermore, Calif.

**July**

Maryann Simmons, UC Berkeley

Matthew Merzbacher, Mills College, Department of Mathematics and Computer Science, Oakland, Calif.

José María Carazo, Univ. Autonoma de Madrid, Madrid, Spain

John Drake, ORNL, Oak Ridge, Tenn.

**August**

Chao Yang, ORNL, Oak Ridge, Tenn.

Patrick Amestoy and Chiara Amestoy, ENSEEIHT-IRIT, Toulouse, France

Jonathan Shewchuk, UC, Berkeley, Computer Science Division, Louis Komzsis, MSC Software, Costa Mesa, Calif.

Gregory T. Balls, UC Berkeley, Computer Science Division

**September**

Lars Jonsson, Department of Physics, Ohio State University Franz-Josef Pfreundt, Institut für Techno- und Wirtschaftsmathematik, Kaiserslautern, Germany

Gunther Weber and Oliver Kreylos, UC Davis, Davis, Calif.

**October**

Daniel C. Glaser, UC Berkeley, School of Information Management and Systems Wah Chiu, Baylor College of Medicine, Biochemistry Department, Houston, Texas

Stephen E. Brenner, UC Berkeley, Department of Plant and Microbial Biology.

**November**

Peter Gottschling, German National Research Center for Information Technology

John Gilbert, Xerox Palo Alto Research Center, Palo Alto, Calif.

**December**

Dan Quinlan, LLNL, Livermore, Calif.

Richard Scalettar, UC Davis, Physics Department, Davis, Calif.

Don Geiser and Tom Morris, API Networks, Concord, Mass.

Hiroaki Matsui, Research Organization for Information Science and Technology, Tokyo, Japan

## STAFF EFFECTIVENESS

*Goal: Leverage staff expertise and capabilities to increase efficiency and effectiveness.*

An ongoing measure of the effectiveness of NERSC's staff is the fact that NERSC is able to offer more computing power with a smaller staff than other computing centers with a comparable budget. Although our staff members certainly do work hard, we think working smarter is what makes the crucial difference. The interdisciplinary expertise of so many staff members makes a major contribution to our effectiveness — their in-depth knowledge of a scientific field as well as computational methods and programming techniques allows them to understand the needs of researchers and make the most effective use of computers to solve scientific problems.

One of our primary objectives is to increase the efficiency and effectiveness of scientific computation, not only for NERSC clients, but within the entire computational science community. This approach is reflected by a user comment to our FY2000 survey, "Very responsive consulting staff that makes the user feel that his problem, and its solution, is important to NERSC." Here are some examples of how we're striving to make our staff as effective as each person can be.

### Recognizing Outstanding Achievement

Several times a year, NERSC honors those employees who have gone beyond their usual job descriptions and responsibilities to make key contributions to the success of our organization. In the year 2000, Outstanding Performance Awards were presented to:

Elizabeth Bautista, for developing and implementing a plan to place all internal documentation on a web site for the Operations Group. The redesigned Web pages make searching for information efficient to the operator, and its central location makes maintenance relatively simple for the webmaster.

Julian Borrill, for his significant contributions to cosmology. By employing the parallel processing power of the Cray T3E supercomputer and by using the MADCAP software package that he developed, he was able to process the 1997 BOOMERanG dataset in just a matter of hours.

Gregory Butler, for leading the effort to conduct testing on the new 608-processor IBM SP system to demonstrate Y2K compliance. Within 10 days after its delivery, Greg had developed a new methodology for Y2K testing on this new architecture, which cut the testing time on the system from the vendor-estimated 4 weeks to 2 weeks.

Brent Draney, for his critical efforts and going to extraordinary lengths in preparing for Y2K (maintaining CUB viability), thereby ensuring that NERSC and Lab operations continued to provide seamless service during the Y2K changeover.

Russell Huie, for taking independent action and placing himself as point of contact for the change in the Video Conference System (VCS). He gathered all the pertinent information needed to operate, troubleshoot and provide same quality of service as required by the ESnet community.

Steven Lowe, for his efforts in all operational aspects of deploying the new HPSS system. He insured that the operations staff were thoroughly familiarized with the new hardware configuration, helped set up the new workstation required for monitoring the system and provided all operations documentation required by the operations staff to efficiently monitor and respond to problems relating to the new storage system.

Leonid Olikier, for his work on programmability, performance and scalability of a mesh adaptation code on a number of parallel architectures. The results, developed jointly with Dr. Rupak Biswas of MRJ Technology Solutions, were reported in a paper titled "Parallelization of a Dynamic Unstructured Applications Using Three Leading Paradigms," which received the "Best Paper of SC99" Award.

Robert Ritchey and Katherine Saucedo, for bringing the NERSC workstations to Y2K compliance. In order to utilize up-to-the-minute patches and to minimize the disruptions to the users, all upgrades were performed after normal business hours beginning on Dec. 16, 1999, and continued through weekends and holidays.

R.K. Owen and John McCarthy for preparing a backup system to provide a needed user account information should NERSC's Central User Bank (CUB) system fail to be operational due to the Y2K rollover. Both took on this extra project on top of their regular duties and had to work prior to the holiday shutdown at an accelerated pace and during the shutdown to meet the unmoving Y2K deadline.

Lynn Rippe, for her efforts as procurement officer responsible for NERSC-3. Lynne succeeded in having IBM deliver a much-improved system that met the requirements.

Dan Gunter, for helping to represent LBNL in the "Grid Performance" working group of the Grid Forum, the new standards organization for emerging Grid technologies.

Marcia Perry, for her initiative in authoring and presenting a paper at the IRMA 2000 conference. Her paper is about her remote camera control and remote videoconferencing control software programs.

Charles McParland, for his work on the AMANDA project — he traveled to the South Pole to install strings of sensors deep in the ice. He was a key part of the design of the sensors and took responsibility for personally helping with the installation.

Gary Hoo, for his outstanding presentation at the Director's review of the NERSC division on April 12 and his special effort as the principal investigator for the QoS:Policy-based Bandwidth Administration FWP.

Inna Dubchak, for her initiative in seeking out scientific interactions with members of the Genome Science Dept. Her contribution resulted in the generation of VISTA, a novel software tool for comparative genomic sequence analysis.

Chris Ding and Yun Helen He, for their activities on climate modeling and simulations at NERSC, including collaborating on two important aspects of climate simulations: input/output (I/O) performance and numerical reproducibility.

Wayne Hurlbert, Stephen Lowe and James Lee, for their perseverance in the face of difficulty and outstanding performance in moving 50 Terabytes of data residing on the IBM libraries to the STK silos. They worked night and day on the hardware failures and also put together a very intensive well-planned effort to move the data.

Nick Cardo, Adrian Wong, James Craw, Richard Gerber, Terri Kaltz, Tina Declerck, Jackie Scoggins, Brent Draney, Mike Welcome, Greg Butler, Jonathon Carter, Tammy Welcome, William Harris, Harsh Anand Passi, Madji Baddourah, Tom DeBoni, Therese Enright, R.K. Owen, Francesca Verdier, William Saphir, Elizabeth Bautista, Martin Stouffer, Andrew Canning, Chris Ding, Esmond Ng, Nancy Meyer, Xiaoye Sherry Li, Osni Marques and David Bailey, for the NERSC-3/Phase 1 effort. The team worked tirelessly for over five months to bring the system into service. In the end, by exceeding the schedule, NERSC was able to provide 1,000,000 more MPP hours than planned at the beginning of the fiscal year.

Marcus Day, for making a major scientific advance in computational science for combustion. He has been a lead developer on a new parallel code for modeling low speed combustion with realistic chemistry.

### **Staff Retention**

NERSC and Berkeley Lab are located near the world's center for high technology. As a result, we must compete with some of the leading names in the field, as well as with a seemingly endless stream of startups. While NERSC cannot offer the stock options and other incentives of competing companies, we can offer a comfortable work environment and opportunities to work on problems of national importance. Overall, we're holding our own in terms of staff retention. In 2000, NERSC experienced a 15 percent turnover rate among staff, which is comparable to private industry.

## PROTECTED INFRASTRUCTURE

*Goal: Provide a secure computing environment for NERSC clients and sponsors.*

Cybersecurity continued to be a hot topic for DOE in 2000, even though the dreaded Y2K debacle failed to materialize. NERSC continued to strengthen its safeguards against outside attacks and to provide a secure computing environment for our users. Although all work conducted using NERSC's systems is unclassified, it is essential that we strike a balance between providing access for users across the country while ensuring that our systems are secure and protected.

BRO, an intrusion detection system developed at Berkeley Lab and named for the ever-watching "Big Brother" in George Orwell's "1984," continued to monitor all incoming and outgoing network traffic to NERSC. BRO actively responds to incoming data packets and searches for unusual patterns which could indicate an attack. If such a pattern is detected, a real-time notification is sent to cyber-security personnel and automatically signals routers to block the suspect traffic.

For the first time, BRO was deployed on an outside system in 2000. Bill Kramer, head of NERSC's High Performance Computing Department, was in charge of building and running the high-speed network serving the SC2000 conference in Dallas and attached BRO to the network to monitor traffic into the Dallas Convention Center.

In 2000, NERSC also created the Networking and Security Group, led by Howard Walter. Because cybersecurity today is almost always related to networking issues, this group was created to develop a unified approach to common problems and solutions. The group regularly reviews cybersecurity issues, ranging from password policies to infrastructure vulnerabilities, investigates attempted attacks and prepares for any potential attacks.

NERSC's vigilance seems to be paying off. Between July 2000 and January 2001, NERSC has contacted other sites regarding security problems with their computers approximately 470 times. During that same period, the number of sites detecting and alerting us about security problems on NERSC systems was zero.

## CONCLUSION

In the arena of high-performance computing, where if you stand still you quickly fall behind, NERSC retains its leadership position, even at the ripe old age of 27. As DOE's flagship computing center for unclassified research, NERSC's success is largely defined by the computational science discoveries of our more than 2,000 users. Essential to this scientific success, however, are the contributions of our staff in providing critical systems and services for our users.

But we also take our definition of success one step farther to include our interaction with HPC vendors, other DOE national laboratories, the scientific community and the HPC community. Successful leadership means working with vendors to improve the ability of computing and storage resources to meet the day-to-day production demands of our users. This experience helps vendors develop better products, while giving other centers the tools to assess how different systems perform in the real world of scientific computing.

NERSC staff members produce hundreds of scientific and technical papers each year, and after careful review, the articles are published for the benefit of other researchers. When organizers of conference need speakers on any number of HPC topics, NERSC staffers are usually on the invitation list. Putting our ideas and experiences out for scrutiny by our peers and having those ideas validated is another measure of our success — and a source of pride for us.

In short, it's all about making a difference. If our work helps a scientist make a significant discovery, helps another facility improve its systems or service, or gives a student a new idea for research, then we've succeeded. With our long track record and string of successes over the years, it's sometimes comforting to pause and look back to see how far our center — and the entire field of HPC — has come. But the real focus of our efforts is on the future and how we can help shape it.